

THE EQUIVALENCE OF TWO SEIBERG-WITTEN FLOER HOMOLOGIES

TYE LIDMAN AND CIPRIAN MANOLESCU

ABSTRACT. We show that monopole Floer homology (as defined by Kronheimer and Mrowka) is isomorphic to the S^1 -equivariant homology of the Seiberg-Witten Floer spectrum constructed by the second author.

CONTENTS

1. Introduction	4
1.1. Background	4
1.2. Results	4
1.3. Outline and organization of the paper	6
1.4. Conventions.	7
1.5. Acknowledgements.	8
2. Morse homology in finite dimensions	9
2.1. The standard construction	9
2.2. Orientations	10
2.3. Quasi-gradient vector fields	13
2.4. Morse homology for isolated invariant sets	17
2.5. Morse homology for manifolds with boundary	18
2.6. Morse homology for manifolds with circle actions	19
2.7. The U -action in Morse homology	23
2.8. Combined generalizations	24
3. The Seiberg-Witten Floer spectrum	26
3.1. The configuration space and the gauge group action	26
3.2. Coulomb slices	26
3.3. The Seiberg-Witten equations	29
3.4. Finite-dimensional approximation	31
3.5. The Conley index and the Seiberg-Witten Floer spectrum	32
4. Monopole Floer homology	34
4.1. Seiberg-Witten equations on the blow-up	34
4.2. Perturbed Seiberg-Witten equations	36
4.3. Tame perturbations	38
4.4. Very tame perturbations	39
4.5. Admissible perturbations	40
4.6. Orientations	43
4.7. Monopole Floer homology	44
4.8. Gradings	46
5. Reduction to the Coulomb gauge	47
5.1. Bundle decompositions and projections	47
5.2. Choices of gauge on cylinders	53
5.3. Controlled Coulomb perturbations	55
5.4. Trajectories in global Coulomb gauge	57
5.5. Hessians	59
5.6. Non-degeneracy of stationary points in Coulomb gauge	69
5.7. Path spaces	71
5.8. Four-dimensional Coulomb slices	73
5.9. The linearized equations	75
5.10. Non-degeneracy of trajectories in Coulomb gauge	81
5.11. Gradings	83
5.12. The cut-down moduli spaces in Coulomb gauge	84
5.13. Orientations	86
5.14. Monopole Floer homology in Coulomb gauge	87

6. Finite-dimensional approximations with tame perturbations	89
6.1. Very compactness	89
6.2. Strategy for the proof of Theorem 1.1	91
6.3. Finite-dimensional approximations of perturbations	92
7. Stationary points	98
7.1. Convergence	98
7.2. Stability	100
7.3. Hyperbolicity	103
7.4. Other stationary points	106
8. The approximate flow as a Morse equivariant quasi-gradient	108
8.1. Control away from the stationary points	109
8.2. The function F_λ	109
8.3. Control in the intermediate region	111
8.4. The L^2 and \tilde{g} metrics	112
8.5. Control near the stationary points	115
8.6. Proposition 8.2 and its consequences	120
9. Gradings	122
9.1. Relative gradings of stationary points	122
9.2. Absolute gradings	130
9.3. Conclusions	133
10. The Morse-Smale condition for the approximate flow	135
11. Self-diffeomorphisms of configuration spaces	137
11.1. The construction of Ξ_λ	139
11.2. Three-dimensional properties of Ξ_λ	141
11.3. Four-dimensional properties of Ξ_λ	143
11.4. Extensions of Ξ_λ to other path spaces	150
12. Convergence of approximate trajectories	153
12.1. Convergence downstairs	153
12.2. Convergence of parameterized trajectories in the blow-up	155
12.3. Near-constant approximate trajectories	158
12.4. Convergence of unparameterized trajectories in the blow-up	162
12.5. Convergence in L_k^2	164
13. Characterization of approximate trajectories	167
13.1. Stability	167
13.2. The U -maps	169
13.3. Orientations	170
14. The equivalence of the homology theories	172
References	176

1. INTRODUCTION

1.1. Background. The Seiberg-Witten (monopole) equations [SW94a, SW94b] are an important tool for understanding the topology of smooth four-dimensional manifolds. A signed count of the solutions of these equations on a closed four-manifold yields the Seiberg-Witten invariant [Wit94]. On a four-manifold with boundary, instead of a numerical invariant one can define an element in a group associated to the boundary, called the Seiberg-Witten Floer homology. There are several different constructions of Seiberg-Witten Floer homology in the literature, [MW01, Man03, KM07, Frø10]. The goal of this paper is to prove that, for rational homology spheres, the definitions given by Kronheimer-Mrowka in [KM07] and by the second author in [Man03] are equivalent.

The construction in [KM07] applies to an arbitrary three-manifold Y , equipped with a Spin^c structure \mathfrak{s} . Given this data, Kronheimer and Mrowka define an infinite dimensional analog of the Morse complex, with the underlying space being the blow-up of the configuration space of Spin^c connections and spinors (modulo gauge). The role of gradient flow lines is played by solutions to generic perturbations of the Seiberg-Witten equations on $\mathbb{R} \times Y$. Their invariant, monopole Floer homology, is the homology of the resulting complex. This complex (and hence also its homology) comes with a $\mathbb{Z}[U]$ -module structure. The applications of monopole Floer homology include the surgery characterization of the unknot [KMOS07] and Taubes' proof of the Weinstein conjecture in three dimensions [Tau07].

In fact, there are three different versions of monopole Floer homology defined in [KM07]; they are denoted \widetilde{HM} , \widehat{HM} , and \overline{HM} . Yet another version, \widetilde{HM} , was constructed by Bloom in [Blo11]: To define \widetilde{HM} , one considers the cone of the U map on the complex that defines \widehat{HM} , and then takes homology.

Compared with [KM07], the construction in [Man03] was originally done only for rational homology spheres; on the other hand, it yields something more than a homology group. By using finite dimensional approximation of the Seiberg-Witten equations, combined with Conley index theory, one obtains an invariant in the form of an equivariant suspension spectrum. Specifically, given a rational homology sphere Y equipped with a Spin^c structure \mathfrak{s} , one can associate to it an S^1 -equivariant spectrum $\text{SWF}(Y, \mathfrak{s})$. (See also [KM02, Kha13, Sas14] for extensions of this construction to the case $b_1 > 0$.)

The S^1 -equivariant homology of $\text{SWF}(Y, \mathfrak{s})$ can be viewed as a definition of Seiberg-Witten Floer homology. The advantage of having a Floer spectrum is that one can also apply other (equivariant) generalized homology functors to it. For example, by adding the conjugation symmetry, one can define a $\text{Pin}(2)$ -equivariant Seiberg-Witten Floer homology; this was instrumental in the disproof of the triangulation conjecture by the second author [Man16]. For other applications of the Floer spectrum, see [Man07, Man14, Lin15, LM14].

1.2. Results. The bulk of this paper is devoted to proving:

Theorem 1.1. *Let Y be a rational homology sphere with a Spin^c structure \mathfrak{s} . There is an isomorphism of absolutely-graded $\mathbb{Z}[U]$ -modules:*

$$\widetilde{HM}_*(Y, \mathfrak{s}) \cong \widetilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s})),$$

where \widetilde{HM} is the “to” version of monopole Floer homology defined in [KM07], and $\widetilde{H}_*^{S^1}$ denotes reduced equivariant (Borel) homology.¹

¹In this paper, we grade Borel homology so that we simply have $\widetilde{H}_*^{S^1}(X) = \widetilde{H}_*(X \wedge_{S^1} ES^1_+)$. This differs from the grading conventions in [GM95] or [KM02] by one.

From Theorem 1.1 we deduce that Bloom's homology \widehat{HM} can be identified with the ordinary (non-equivariant) homology of the Floer spectrum SWF:

Corollary 1.2. *Let Y be a rational homology sphere equipped with a Spin^c structure \mathfrak{s} . Then, $\widehat{HM}_*(Y, \mathfrak{s}) \cong \widetilde{H}_*(\text{SWF}(Y, \mathfrak{s}))$ as absolutely graded abelian groups.*

From the absolute grading on monopole Floer homology one can extract a \mathbb{Q} -valued invariant, called the Frøyshov invariant; see [Frø10] or [KM07, Section 39.1]. A similar numerical invariant, called δ , was defined in [Man16, Section 3.7] using the Floer spectrum SWF. (The definition there was only given for Spin structures, but it extends to the Spin^c setting.) An immediate consequence of Theorem 1.1 is

Corollary 1.3. *Let Y be a rational homology sphere equipped with a Spin^c structure \mathfrak{s} . Then, $\delta(Y, \mathfrak{s}) = -h(Y, \mathfrak{s})$, where h is the Frøyshov invariant as defined in [KM07, Section 39.1].*

Furthermore, we can combine Theorem 1.1 with the equivalence between monopole Floer homology and Heegaard Floer homology, established in work of Kutluhan-Lee-Taubes [KLT10a, KLT10b, KLT10c, KLT11, KLT12] and of Colin-Ghiggini-Honda [CGH12b, CGH12c, CGH12a] and Taubes [Tau10]. In this way we obtain a relationship between the spectrum $\text{SWF}(Y, \mathfrak{s})$ and Heegaard Floer theory. Precisely, we confirm a conjecture from [KM02], that the different flavors of Heegaard Floer homology are different equivariant homology theories applied to $\text{SWF}(Y, \mathfrak{s})$:

Corollary 1.4. *If Y is a rational homology sphere and \mathfrak{s} is a Spin^c structure on Y , we have the following isomorphisms of relatively graded $\mathbb{Z}[U]$ -modules:*

$$\begin{aligned} HF_*(Y, \mathfrak{s})^+ &\cong \widetilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s})), & \widehat{HF}_*(Y, \mathfrak{s}) &\cong \widetilde{H}_*(\text{SWF}(Y, \mathfrak{s})), \\ HF_*(Y, \mathfrak{s})^- &\cong c\widetilde{H}_*(\text{SWF}(Y, \mathfrak{s})), & HF_*^\infty(Y, \mathfrak{s}) &\cong t\widetilde{H}_*(\text{SWF}(Y, \mathfrak{s})), \end{aligned}$$

where $c\widetilde{H}^*$ and $t\widetilde{H}^*$ denote co-Borel and Tate homology, respectively, as defined in [GM95].

The results above can be applied in two directions. On the one hand, there are numerous computational techniques available for Heegaard Floer homology, whereas the class of manifolds Y for which one can compute $\text{SWF}(Y, \mathfrak{s})$ is rather small. By making use of the isomorphisms in Corollary 1.4, one can at least understand the (equivariant or non-equivariant) homologies of $\text{SWF}(Y, \mathfrak{s})$. In particular:

- If $\widehat{HF}_*(Y, \mathfrak{s}) \cong \mathbb{Z}$, this suffices to determine the non-equivariant stable homotopy type of $\text{SWF}(Y, \mathfrak{s})$: by the Hurewicz and Whitehead theorems, it has to be that of a (de-)suspension of the sphere spectrum. (However, it is unclear whether the equivariant stable homotopy type of $\text{SWF}(Y, \mathfrak{s})$ is determined by this information.)
- In the case when \mathfrak{s} is a Spin structure, the Seiberg-Witten equations on (Y, \mathfrak{s}) have a $\text{Pin}(2)$ symmetry, and one can define the $\text{Pin}(2)$ -equivariant homology of $\text{SWF}(Y, \mathfrak{s})$. (This homology was used in [Man16] to disprove the triangulation conjecture in high dimensions.) When Y is Seifert fibered, the Heegaard Floer homology was calculated in [OS03], so Corollary 1.4 can tell the S^1 -equivariant homology of $\text{SWF}(Y, \mathfrak{s})$. Together with knowledge of the monopoles (generators of the monopole Floer complex) from [MOY97], this suffices to determine the $\text{Pin}(2)$ -equivariant homology of $\text{SWF}(Y, \mathfrak{s})$; see [Sto] for details.

In the reverse direction, there are certain results that are easier to prove with Floer spectra, and one can use Theorem 1.1 and Corollary 1.4 to translate them into the settings of monopole Floer or Heegaard Floer homology. This is the case with the Smith inequality for the Floer homology of coverings, which is the subject of the sequel to this paper [LM14]. By combining the Smith inequality with the knot surgery formula in Heegaard Floer homology [OS11], we obtain restrictions on surgeries on a knot being regular covers over other surgeries on that knot, and over surgeries on other knots; see [LM14].

1.3. Outline and organization of the paper. Theorem 1.1 asserts that the monopole Floer homology of Kronheimer-Mrowka is isomorphic to the equivariant homology of the Conley index used for $\text{SWF}(Y, \mathfrak{s})$. Let us describe the basic strategy for the proof. The monopole Floer complex is defined from the *perturbed* Seiberg-Witten equations in infinite dimensions, whereas the Conley index is a space associated to an *approximation* of the Seiberg-Witten equations in finite dimensions. On the Conley index side, we can also use an approximation of the perturbed Seiberg-Witten equations; by the continuation properties of the Conley index, this yields a space homotopy equivalent to the one from the unperturbed case. Thus, we start by fixing a suitable perturbation, so that we have regularity in infinite dimensions, and hence we can define the monopole Floer complex. We then show that for a large enough approximation, we also have regularity in finite dimensions, and hence the homology of the Conley index is given by a Morse complex. Further, by applying the inverse function theorem (for stationary points and trajectories), we conclude that the Morse complex in finite dimensions is isomorphic to the original monopole Floer complex, and this will complete the proof.

There are several technical difficulties that must be overcome to implement this strategy. To start with, the version of Morse homology that we need to use in finite dimensions is different from the standard one in three ways: we are on a non-compact manifold (a vector space), we have to define S^1 -equivariant rather than ordinary homology, and the approximate Seiberg-Witten flow is not a gradient flow. Non-compactness is taken care of using the notion of Conley index; this goes back to the work of Floer [Flo89]. Equivariance was dealt with by Kronheimer and Mrowka in [KM07], using the real blow-up construction. One intriguing aspect is that the flow is not a gradient. This leads us to introduce the weaker notion of *quasi-gradient*, which is a particular case of the Morse-Smale flows that appear in the study of stability for dynamical systems. We will show that one can define Morse homology from quasi-gradients. All this is done in Section 2; there, we recall the various notions from Morse theory and Conley index theory, the relation between the two theories, and then proceed to define (in several steps) equivariant Morse homology for quasi-gradients on non-compact manifolds.

In Section 3 we outline the construction of the Seiberg-Witten Floer spectrum from [Man03]. The main idea is to approximate the Seiberg-Witten equations (in global Coulomb gauge) by a flow in finite dimensions, and then to take the Conley index associated to that flow.

In Section 4 we sketch the construction of monopole Floer homology by Kronheimer and Mrowka, following their book [KM07]. In particular, we describe the class of admissible perturbations of the Seiberg-Witten equations that can be used to define the Floer complexes.

In Section 5 we explain how the Kronheimer-Mrowka construction can be rephrased in terms of configurations in global Coulomb gauge. This is the first step in bringing it closer

to the construction of the Floer spectrum. The section is rather lengthy, because there are several aspects of the theory that have to be translated into the new setting: gauge conditions (in particular, for four-dimensional configurations we introduce the concept of *pseudo-temporal gauge*), admissibility for perturbations, Hessians, regularity for stationary points, spaces of paths, regularity for trajectories, gradings, orientations, and the $\mathbb{Z}[U]$ -module action.

Another step in relating the two theories is taken in Section 6. There, we show that the Floer spectrum can be defined from finite dimensional approximation of the *perturbed* Seiberg-Witten equations. We use an admissible perturbation of the kind considered in [KM07].

In Section 7 we use the inverse function theorem to identify the stationary points in infinite dimensions (i.e., the generators of the monopole Floer complex) with the stationary points of the gradient flow in the finite dimensional approximations. (This identification is limited to a certain grading range.) Further, we show that if the stationary points are non-degenerate in infinite dimensions, then the corresponding stationary points in sufficiently large approximations are non-degenerate as well. Note that the approximations are described by eigenvalue cut-offs, and these range over the interval $(0, \infty]$, with ∞ giving the original equations (in infinite dimensions). When applying the inverse function theorem, one thing to be careful about is how to give a manifold-with-boundary structure to the interval $(0, \infty]$. We do so by identifying $(0, \infty]$ with $[0, 1]$, via a homeomorphism that depends on the growth rate of the eigenvalues.

In Section 8 we show that the approximate Seiberg-Witten flow in finite dimensions is a quasi-gradient. The main difficulty there is to perturb the Chern-Simons-Dirac functional in such a way so that it decreases along flow lines.

In Section 9, we study the gradings of the stationary points in the finite dimensional approximations and show that the correspondence with stationary points in infinite dimensions established in Section 7 preserves gradings.

In Section 10, we show how to arrange so that the Morse-Smale condition is satisfied for the approximate flow.

Section 11 contains the construction of some diffeomorphisms Ξ_λ of the configuration space. These diffeomorphisms take a stationary point of the approximate Seiberg-Witten flow to the corresponding stationary point of the original flow. They allow us to identify the corresponding path spaces, so that we may directly relate infinite dimensional trajectories with approximate trajectories in Sections 12 and 13.

In Section 12 we prove various convergence results for trajectories of the approximate Seiberg-Witten flow, in the limit as the eigenvalue cut-off goes to infinity.

In Section 13 we use the inverse function theorem to identify the flow trajectories in infinite dimensions with those in the approximations. Furthermore, we show that these identifications preserve orientations.

Finally, in Section 14, we bring these results together to prove Theorem 1.1, about the equivalence between monopole Floer homology and the equivariant homology of $\text{SWF}(Y, \mathfrak{s})$. We also prove the three corollaries stated in Section 1.2.

1.4. Conventions. Throughout the paper, when we talk about the flow associated to a vector field v , we mean the reverse flow generated by v , i.e, the flow whose trajectories satisfy:

$$\frac{d}{dt}\gamma(t) + v(\gamma(t)) = 0.$$

Also, we depart from the terminology in [KM07], where critical points of the CSD functional are also referred to as critical points of the Seiberg-Witten vector field. Here, we sometimes need to think of vector fields as maps to the tangent space, and the critical points of a map are the zeros of its derivative. To prevent confusion, we will use the term *stationary point* to refer to the zero of a vector field.

1.5. Acknowledgements. We thank Peter Kronheimer, Robert Lipshitz, Max Lipyanskiy, Tim Perutz, and Matthew Stoffregen for helpful conversations.

2. MORSE HOMOLOGY IN FINITE DIMENSIONS

In this section we describe several versions of Morse homology (all in finite dimensions), building up to the construction that will be needed in the paper.

2.1. The standard construction. Let X be a closed, oriented Riemannian manifold equipped with a Morse function f . The Morse-Smale condition requires that for any two critical points x and y , the unstable manifold W_x^u and the stable manifold W_y^s intersect transversely. Their intersection is then a smooth manifold $M(x, y)$, the moduli space of parameterized flows between x and y , which has dimension $\text{ind}(x) - \text{ind}(y)$.

With this information, we can build the *Morse complex* [Wit82, Bot88, Flo89]. The chain groups in degree i , denoted C_i , are freely generated over \mathbb{Z} by the critical points of f of index i . We define the differential ∂ as follows. If x and y are critical points of index i and $i - 1$ respectively, then $\check{M}(x, y) := M(x, y)/\mathbb{R}$ is a compact, oriented 0-manifold. (See Section 2.2 below for the construction of orientations.) We let $n(x, y)$ denote the number of unparameterized flow lines of ∇f from x to y , counted with sign. The Morse differential is

$$(1) \quad \partial x = \sum_{\{y | \text{ind}(y) = \text{ind}(x) - 1\}} n(x, y) \cdot y.$$

It turns out that $\partial^2 = 0$ and one can take the homology of this complex. The Morse homology is isomorphic to the singular homology, $H_*(X)$.

For future reference, we mention a (well-known) alternative way of expressing the Morse-Smale condition. Let $f : M \rightarrow \mathbb{R}$ be Morse, and x, y be two critical points. Define $\mathcal{P}(x, y)$ to be the space of smooth paths $\gamma : \mathbb{R} \rightarrow M$ with $\lim_{t \rightarrow -\infty} \gamma(t) = x$ and $\lim_{t \rightarrow +\infty} \gamma(t) = y$. The space $\mathcal{P}(x, y)$ has a well-defined L_k^2 completion $\mathcal{P}_k(x, y)$, which is a Banach manifold. (See, for example, [Sch93, Section 2.1] for the case $k = 1$.) For $0 \leq j \leq k$, we let $T_{j, \gamma} \mathcal{P}(x, y)$ be the L_j^2 completion of the tangent space to $\mathcal{P}_k(x, y)$ at γ . (If we fix j , these spaces are naturally identified for all $k \geq j$, so we do not include k in the notation.) The moduli space $M(x, y)$ is the zero set of the section

$$F : \mathcal{P}_k(x, y) \rightarrow T_{k-1} \mathcal{P}(x, y), \quad \gamma \mapsto \frac{d\gamma}{dt} + (\nabla f)(\gamma).$$

Given a path $\gamma \in \mathcal{P}_k(x, y)$, we get the linearized operator

$$(2) \quad L_\gamma = (dF)_\gamma : T_{j, \gamma} \mathcal{P}(x, y) \rightarrow T_{j-1, \gamma} \mathcal{P}(x, y),$$

for all j with $1 \leq j \leq k$.

The definition of L_γ involves taking derivatives of vector fields, i.e., covariant derivatives. To do this, we need to choose a connection on TM in a neighborhood U of the image of γ . This can be the Levi-Civita connection coming from the Riemannian metric or, as in [AD14], the trivial connection induced from a trivialization of TM over U . In either case, if we denote by D the covariant derivative, and $w \in T_{j, \gamma} \mathcal{P}(x, y)$ is a vector field along γ , we have

$$(3) \quad L_\gamma(w) = \frac{Dw}{dt} + D(\nabla f)(w).$$

Note that if D is the Levi-Civita connection, then the second term $D(\nabla f)(w)$ is the Hessian of f applied to w at $\gamma(t)$.

One can check that L_γ is a Fredholm operator. When $\gamma \in M(x, y)$, the index of L_γ is the expected dimension of the moduli space $M(x, y)$.

Lemma 2.1. *Suppose $k \geq 2$ and $1 \leq j \leq k$. Let D and D' be two connections on TM in an open set U containing the image of $\gamma \in \mathcal{P}_k(x, y)$. Write L_γ and L'_γ for the operators (2) that correspond to D , resp. D' . Then:*

- (i) *The difference $L_\gamma - L'_\gamma$ is compact, so the operators L_γ and L'_γ have the same Fredholm index;*
- (ii) *If $\gamma \in M(x, y)$, then $L_\gamma = L'_\gamma$.*

Proof. Let $A = D - D' \in \Omega^1(U; \text{End}(TM))$. Then

$$(4) \quad L_\gamma - L'_\gamma = A\left(\frac{d\gamma}{dt} + (\nabla f)(\gamma)\right).$$

Since $k \geq 2$ and $j \leq k$, we have a Sobolev multiplication $L^2_{j-1} \times L^2_{k-1} \rightarrow L^2_{j-1}$. From here we see that the difference $L_\gamma - L'_\gamma$ takes $T_{j-1, \gamma}\mathcal{P}(x, y)$ to itself.

Consider a sequence of smooth bump functions $\beta_n : \mathbb{R} \rightarrow [0, 1]$ such that β_n is supported on $[-n-1, n+1]$ and identically 1 on $[-n, n]$. The truncations

$$\beta_n \cdot (L_\gamma - L'_\gamma) : T_{j, \gamma}\mathcal{P}(x, y) \rightarrow T_{j, \gamma}\mathcal{P}(x, y)$$

converge to $L_\gamma - L'_\gamma$ in operator norm. Furthermore, when precomposing these truncations with the inclusion of $T_{j, \gamma}\mathcal{P}(x, y)$ into $T_{j-1, \gamma}\mathcal{P}(x, y)$, we obtain compact operators. (Indeed, we can apply Rellich's lemma, because we restrict attention to a fixed compact interval.) Since the limit of compact operators is compact, we conclude that $L_\gamma - L'_\gamma$, as an operator from $T_{j, \gamma}\mathcal{P}(x, y)$ to $T_{j-1, \gamma}\mathcal{P}(x, y)$, is compact.

Part (ii) follows immediately from (4). \square

Lemma 2.2. *Fix $j \geq 1$. Then, the function f is Morse-Smale if and only if for any two critical points x and y of f , and for any gradient trajectory $\gamma \in M(x, y)$, the operator L_γ from (2) is surjective.*

Proof. See [AD14, Theorem 10.1.5] for the case $j = 1$. The main idea is that surjectivity of L_γ is equivalent to injectivity of the formal adjoint L_γ^* and, for any $t \in \mathbb{R}$, one can identify $\ker(L_\gamma^*)$ with the orthogonal complement $(T_{\gamma(t)}W_x^u + T_{\gamma(t)}W_y^s)^\perp$. Thus, W_x^u and W_y^s are transverse at $\gamma(t)$ if and only if L_γ^* is injective.

The same proof works for any $j \geq 1$. In fact, the kernel of $L_\gamma^* : T_{j, \gamma}\mathcal{P}(x, y) \rightarrow T_{j-1, \gamma}\mathcal{P}(x, y)$ consists of smooth configurations (as can be seen by a standard bootstrapping argument), and hence the kernel is independent of j . \square

2.2. Orientations. We now discuss the construction of orientations in Morse homology. The traditional way is to choose orientations for W_x^u for all critical points x . These induce orientations on W_x^s , and hence on the moduli spaces $M(x, y)$. When $\text{ind}(x) - \text{ind}(y) = 1$, a trajectory $u \in M(x, y)$ is counted in the differential ∂ with a sign ± 1 depending on whether the orientation of $M(x, y)$ at u coincides with the canonical orientation of \mathbb{R} .

An alternative way of constructing orientations, closer to Floer theory, is described in [Sch93, Chapter 3]. Let us sketch this construction here. First, note that to every Fredholm operator L we can associate a one-dimensional vector space, the determinant line

$$\det(L) = \Lambda^{\max}(\ker L) \otimes \Lambda^{\max}(\text{coker } L)^*.$$

From the proof of Lemma 2.2 we see that the tangent bundle to $M(x, y)$ can be identified with the kernel of the operator L_γ from (3), whereas the cokernel of L_γ is trivial. Thus, orienting $M(x, y)$ is equivalent to orienting the determinant line bundle $\det(L) \rightarrow M(x, y)$, with fibers $\det(L_\gamma)$ over $\gamma \in M(x, y)$.

Following [Sch93, Definition 2.1], let us equip $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$ with the structure of a smooth manifold with boundary by requiring that the map $\overline{\mathbb{R}} \rightarrow [-1, 1], t \mapsto t/\sqrt{1+t^2}$ is a diffeomorphism. For every $x, y \in X$ (not necessarily critical points), we consider the space of smooth curves

$$C_{x,y}^\infty = \{\gamma \in C^\infty(\overline{\mathbb{R}}, X) \mid \gamma(-\infty) = x, \gamma(+\infty) = y\}.$$

For $\gamma \in C_{x,y}^\infty$, we let Σ_{γ^*TX} be the space of Fredholm operators of the form

$$(5) \quad K = \frac{D}{dt} + A(t) : T_{1,\gamma}\mathcal{P}(x, y) \rightarrow T_{0,\gamma}\mathcal{P}(x, y),$$

where $A \in C^0(\mathbb{R}; \text{End}(\gamma^*TX))$ is a path of endomorphisms such that $K^\pm := A(\pm\infty)$ is non-degenerate and conjugated self-adjoint². In particular, when x and y are critical points, the operator L_γ from (3) is of this form.

Consider the set of pairs (γ, K) , where $\gamma \in C^\infty(\overline{\mathbb{R}}, X)$ and $K \in \Sigma_{\gamma^*TX}$. Following [Sch93, Definition 3.7], two such pairs (γ, K) and (ζ, L) are said to be equivalent if we have the asymptotical identities

$$\gamma(\pm\infty) = \zeta(\pm\infty) \quad \text{and} \quad K^\pm = L^\pm.$$

Let \mathcal{E} be the set of such equivalence classes. Thus, an element of \mathcal{E} is determined by two points $x, y \in X$, together with two non-degenerate and conjugated self-adjoint operators $K^- \in \text{End}(T_x X)$ and $K^+ \in \text{End}(T_y Y)$.

It is proved in [Sch93, Section 3.2.1] that, if (γ, K) is equivalent to (ζ, L) , there is a natural identification of the determinant lines $\det(K)$ and $\det(L)$. Hence, we can write $\det([\gamma, K])$ for a class $[\gamma, K] \in \mathcal{E}$. As in [Sch93, Definition 3.15], we define a *coherent orientation* to be a map σ that associates an orientation of $\det([\gamma, K])$ to any $[\gamma, K] \in \mathcal{E}$, in a way compatible with concatenation of paths and operators; that is,

$$(6) \quad \sigma[\gamma, K] \# \sigma[\zeta, L] = \sigma[\gamma \# \zeta, K \# L],$$

whenever $\gamma(+\infty) = \zeta(-\infty)$ and $K^+ = L^-$. We denote by $\#$ the natural concatenation operator.

Proposition 3.16 in [Sch93] shows that coherent orientations exist. Concretely, a coherent orientation can be specified by choosing a basepoint $x_0 \in X$, some non-degenerate and conjugated self-adjoint operator $A \in \text{End}(T_{x_0} X)$, choosing the trivial orientation for the class $[x_0, K_0]$, where x_0 is the constant path at x_0 and $K_0 = \frac{d}{dt} + A$, and finally choosing arbitrary orientations $\sigma[\gamma, K]$ for all classes $[\gamma, K] \in \mathcal{E}$ with $[\gamma, K] \neq [x_0, K_0]$ and $\gamma(-\infty) = x_0$. Once this set of data is fixed, the orientations for the other classes in \mathcal{E} are determined by the gluing condition (6).

In particular, a coherent orientation gives orientations of the determinant line bundles $\det(L) \rightarrow M(x, y)$, and hence of the moduli spaces $M(x, y)$ themselves. These orientations can be used to define the Morse differential as in (1). Furthermore, it is proved in [Sch93, Appendix B] that the coherent orientation can be chosen so that the orientations on $M(x, y)$ coincide with those coming from orienting the unstable manifolds. Therefore, the resulting Morse complex is the same as in the classical definition.

We now introduce a third way of defining orientations for the Morse complex. This is a variant of Schwarz's construction, but closer to what Kronheimer and Mrowka do for

²In [Sch93], for a real vector space V , an operator $A \in \text{End}(V)$ is called conjugated self-adjoint if it is self-adjoint with respect to some scalar product. This condition is equivalent to A being diagonalizable over \mathbb{R} .

monopole Floer homology in [KM07, Section 20]. Rather than considering all Fredholm operators of the form $\frac{D}{dt} + A(t)$, we focus on the operators L_γ as in (3), but with $x, y \in X$ and $\gamma \in \mathcal{P}(x, y)$ arbitrary. However, note that L_γ is only Fredholm when the Hessian $\text{Hess}(f) = D(\nabla f)$ is non-degenerate at its endpoints x, y . This is not true in general, so let us instead consider compact intervals $I = [t_1, t_2] \subset \mathbb{R}$, spaces of paths

$$\mathcal{P}^I(x, y) = \{\gamma \in C^\infty(I, X) \mid \gamma(t_1) = x, \gamma(t_2) = y\},$$

and their L_k^2 completions $\mathcal{P}_k^I(x, y)$. At each $x \in X$, the Hessian $\text{Hess}(f)_x$ is self-adjoint, and gives a decomposition of $T_x X$ into the span of the nonpositive and positive eigenspaces:

$$T_x X = H_x^- \oplus H_x^+.$$

In other words, we are using the spectral decomposition of the non-degenerate self-adjoint operator $\text{Hess}(f)_x - \epsilon$, for $\epsilon > 0$ small.

Let Π_x^- and Π_x^+ be the orthogonal projections onto H_x^- and H_x^+ , respectively. For any $\gamma \in \mathcal{P}_k^I(x, y)$, we define a Fredholm operator

$$(7) \quad \begin{aligned} \tilde{L}_\gamma : T_{1,\gamma} \mathcal{P}^I(x, y) &\rightarrow T_{0,\gamma} \mathcal{P}^I(x, y) \oplus H_x^+ \oplus H_y^- \\ w &\mapsto \left(\frac{Dw}{dt} + \text{Hess}(f)(w), -\Pi_x^+(w(t_1)), \Pi_y^+(w(t_2)) \right). \end{aligned}$$

One can turn \tilde{L}_γ into an operator of the form (5) as follows. Pick a smooth map $f : I \rightarrow I$ such that $f' \geq 0$ and

$$f(t_1 + s) = t_1, \quad f(t_2 - s) = t_2$$

for all $s \in [0, \delta]$, for some small $\delta > 0$. Then, consider the extended path

$$\gamma^{\text{ext}} : \mathbb{R} \rightarrow X, \quad \gamma^{\text{ext}}(t) = \begin{cases} x & \text{for } t < t_1, \\ \gamma(f(t)) & \text{for } t \in I, \\ y & \text{for } t > t_2. \end{cases}$$

Define a Fredholm operator $\tilde{L}_\gamma^{\text{ext}} \in \Sigma_{(\gamma^{\text{ext}})^* TX}$ by

$$\tilde{L}_\gamma^{\text{ext}} = \frac{D}{dt} + \text{Hess}(f)_{\gamma^{\text{ext}}(t)} - \epsilon,$$

for $\epsilon > 0$ small. Standard deformation and concatenation arguments show that the operators \tilde{L}_γ and $\tilde{L}_\gamma^{\text{ext}}$ have the same index, and that orienting $\det(\tilde{L}_\gamma)$ is equivalent to orienting $\det(\tilde{L}_\gamma^{\text{ext}})$.

Let $\Lambda_\gamma(x, y)$ be the two-element set consisting of the orientations of $\det(\tilde{L}_\gamma)$ or, equivalently, of $\det(\tilde{L}_\gamma^{\text{ext}})$. For fixed x, y but different $\gamma \in \mathcal{P}_k(x, y)$, the pairs $(\gamma^{\text{ext}}, \tilde{L}_\gamma^{\text{ext}}) \in \mathcal{E}$ are equivalent. Hence, by the results of [Sch93, Section 3.2.1], we have a canonical identification between the different sets $\Lambda_\gamma(x, y)$. Consequently, we can drop γ from the notation and write $\Lambda(x, y)$ for any $\Lambda_\gamma(x, y)$.

By considering concatenation of paths, we obtain a natural composition map

$$(8) \quad \Lambda(x, y) \times \Lambda(y, z) \rightarrow \Lambda(x, z).$$

Definition 2.3. A specialized coherent orientation o consists of choices of elements $o_{x,y} \in \Lambda(x, y)$, one for each $x, y \in X$, such that

$$o_{x,y} \cdot o_{y,z} = o_{x,z}$$

for all $x, y, z \in X$, where the multiplication is with respect to (8).

A specialized coherent orientation can be constructed as follows. We choose a critical point x_0 of f , and consider the operator \tilde{L}_{x_0} for the constant path at x_0 . There is a canonical identification $\det(\tilde{L}_{x_0}) \cong \mathbb{R}$, and hence from $o_{x_0, x_0} \cong \{\pm 1\}$ we can choose the element $+1$. Then, we choose arbitrary elements of $o_{x_0, x}$ for all other $x \in X$, and we obtain elements in any $o_{x, y}$ using the concatenation rule.

Using the identification between $\det(\tilde{L}_\gamma)$ and $\det(\tilde{L}_\gamma^{\text{ext}})$, we see that a coherent orientation gives rise to a specialized coherent orientation. Conversely, any specialized coherent orientation can be extended to a coherent orientation. To check this last statement, consider the construction of a specialized coherent orientation from the choices of x_0 and elements in $o_{x_0, x}$, as above. This is a subset of the data needed to give a coherent orientation. Indeed, letting $K_0 = \frac{d}{dt} + \text{Hess}(f)_{x_0}$, we see that from the elements of $o_{x_0, x}$ we get orientations $\sigma[\gamma, K]$ for paths γ starting at x_0 and ending at x , and for K with $K^+ = H_x^+$. Let us now also choose orientations $\sigma[\gamma, K]$ for $\gamma \in \mathcal{P}_k(x_0, x)$ but $K^+ \neq H_x^+$. This gives the necessary data to construct a coherent orientation.

Finally, note that a specialized coherent orientation gives rise to orientations on the moduli spaces $M(x, y)$ in the following way. Given $\gamma \in M(x, y)$, pick $I = [t_1, t_2] \subset \mathbb{R}$ large enough so that the operators $\text{Hess}(f)_{\gamma(t)}$ are non-degenerate for all $t \in (-\infty, t_1] \cup [t_2, \infty)$. Consider the restriction of γ to the interval I , and the corresponding operator $\tilde{L}_{\gamma|_I}$. The orientation on $\det(\tilde{L}_{\gamma|_I})$ gives an orientation on $\det(L_\gamma)$, by a concatenation process similar to the one in [KM07, Section 20.4]. In turn, this orients the tangent space $T_\gamma M(x, y)$.

We conclude that a specialized coherent orientation o can be used to give signs in the Morse complex. When o is induced by a coherent orientation σ , the signs coming from o are the same as those coming from σ . Thus, the three ways of producing orientations, which we have described in this section, all produce the same Morse complex.

2.3. Quasi-gradient vector fields. The purpose of this subsection is to extend Morse-Smale theory to a class of vector fields that are not gradients. Note that among all smooth vector fields on a manifold, gradient vector fields are rather special—even when allowing both the metric and the function to vary. For example, the differential of a gradient vector field at a critical point has only real eigenvalues (because it is a symmetric operator with respect to the given metric). This is also true for the gradient-like (also called pseudo-gradient) vector fields that are sometimes used in Morse theory, e.g., in [Mil65, AD14]. We want to relax this condition, and allow for the stationary points of our vector field to only be hyperbolic, in the following sense:

Definition 2.4. *Let v be a smooth vector field on a manifold X . A stationary point x of v is called hyperbolic if (the complexification of) the derivative $(dv)_x : T_x X \rightarrow T_x X$ has no purely imaginary eigenvalues.*

We now introduce the class of vector fields we want to work with.

Definition 2.5. *Let X be a smooth manifold. A smooth vector field v on X is called Morse quasi-gradient if the following conditions are satisfied:*

- (a) *All stationary points of v are hyperbolic.*
- (b) *There exists a smooth function $f : X \rightarrow \mathbb{R}$ such that $df(v) \geq 0$ at all $x \in X$, with equality holding if and only if x is a stationary point of v .*

Example 2.6. The gradient vector field of a Morse function f has hyperbolic stationary points, since the eigenvalues of $(dv)_x$ are real and non-zero; further, the Morse function f can be used in condition (b).

For future reference, it is worth pointing out the following.

Lemma 2.7. *Let v be a Morse quasi-gradient vector field on a smooth manifold X , and let f be a function as in part (b) of Definition 2.5. Then, any stationary point of v is a critical point of f .*

Proof. Let x be a stationary point of v . Since x is a minimum of the function $df(v)$, we have $d(df(v))_x = 0$. For $y \in T_x X$, we get

$$0 = d(df(v))_x(y) = (d^2 f)_x(v_x, y) + (df)_x((dv)_x(y)).$$

The first term in the last expression is zero, because $v_x = 0$. Further, since x is hyperbolic, we have that $(dv)_x$ is an automorphism of $T_x X$, hence $(dv)_x(y) = z$ can take any value in $T_x X$. We conclude that $(df)_x(z) = 0$ for all $z \in T_x X$, so x is a critical point of f . \square

For the rest of this subsection we will assume that X is closed and oriented. Let v be a Morse quasi-gradient vector field on X . We seek to establish several properties that v has in common with gradients of Morse functions.

Let \mathfrak{C} be the set of stationary points of v . If $x \in \mathfrak{C}$, hyperbolicity implies that we have a decomposition

$$T_x X = T_x^s X \oplus T_x^u X,$$

where $T_x^s X$ and $T_x^u X$ are invariant subspaces for $L = (dv)_x$, the eigenvalues of $L|_{T_x^s X}$ have negative real part, and the eigenvalues of $L|_{T_x^u X}$ have positive real part. (See [PdM82, Proposition 2.8] for a proof.) We define the *index* of x to be the dimension of $T_x^u X$.

The local structure of flows near hyperbolic singularities is well-understood; a good reference is [PdM82, Chapter 2]. Theorem 4.10 in [PdM82] says that around every $x \in \mathfrak{C}$ there is a local coordinate chart \mathcal{N}_x such that the restriction of v to \mathcal{N}_x is conjugate to a linear flow (under a homeomorphism). This implies that x is the only stationary point inside \mathcal{N}_x . Since X is compact, we deduce that the set \mathfrak{C} is finite.

We can arrange so that the neighborhoods \mathcal{N}_x are disjoint for different x . By part (b) of Definition 2.5, there is an $\epsilon > 0$ such that

$$(9) \quad df(v)(p) \geq \epsilon, \quad \forall p \in X - \bigcup_{x \in \mathfrak{C}} \mathcal{N}_x.$$

Next, consider the flow on X associated to v . Let $\gamma : \mathbb{R} \rightarrow X$ be a flow line, i.e., $d\gamma/dt + v(\gamma(t)) = 0$. We have:

$$(10) \quad \frac{d}{dt} f(\gamma(t)) = -df_{\gamma(t)} \circ v(\gamma(t)) \leq 0,$$

with equality only if $\gamma(t)$ is stationary. Therefore, f decreases (strictly) along non-stationary flow lines.

Lemma 2.8. *Let v be a Morse quasi-gradient vector field on a closed, oriented, smooth manifold X . Let $\gamma : \mathbb{R} \rightarrow X$ be a flow line of v . Then, $\lim_{t \rightarrow -\infty} \gamma(t)$ and $\lim_{t \rightarrow +\infty} \gamma(t)$ exist, and they are both stationary points of v .*

Proof. This is similar to the Morse gradient case; cf. [BH04, Proposition 3.19]. Let \mathcal{N}_x be the neighborhoods of $x \in \mathfrak{C}$ constructed above, and $\epsilon > 0$ be such that (9) holds.

Compactness of X implies that $\gamma(t)$ is defined for all $t \in \mathbb{R}$, and $f \circ \gamma : \mathbb{R} \rightarrow \mathbb{R}$ is bounded. (Here, f is the function in part (b) of Definition 2.5.) From Equation (10) we see that

$$\lim_{t \rightarrow \pm\infty} -df(v)(\gamma(t)) = \lim_{t \rightarrow \pm\infty} \frac{d}{dt} f(\gamma(t)) = 0.$$

In view of (9), this shows that for $t \ll 0$, the flow line $\gamma(t)$ is contained in the chart \mathcal{N}_x for some x . Further, if $t_n \rightarrow -\infty$, then at any accumulation point of the sequence $\{\gamma(t_n)\} \subseteq X$ we must have $df(v) = 0$. Hence, the only accumulation point (which exists by the compactness of X) must be the stationary point x . We deduce that $\lim_{t \rightarrow -\infty} \gamma(t) = x$. A similar argument applies to $\gamma(t)$ as $t \rightarrow +\infty$. \square

Hyperbolicity of stationary points implies the existence of stable and unstable manifolds W_x^s, W_x^u , cf. the Stable Manifold Theorem [PdM82, Theorem 6.2]. These are the images of injective immersions

$$\begin{aligned} E^s : T_x^s X &\rightarrow W_x^s \subseteq X, \\ E^u : T_x^u X &\rightarrow W_x^u \subseteq X. \end{aligned}$$

If we did not have condition (b) in Definition 2.5, the maps E^s and E^u may not be embeddings, and W_x^s, W_x^u may not be manifolds. (See [PdM82, p.74] for an example.) However, for Morse quasi-gradients, E^s and E^u are homeomorphisms onto their images, and hence embeddings. The proof of this fact can be taken verbatim from the Morse gradient case; cf. Lemma 4.10 in [BH04]. Indeed, the main input in that proof is the existence of limits of flow lines at $\pm\infty$; this was established for quasi-gradients in Lemma 2.8.

Thus, for every $x \in \mathfrak{C}$, we have stable and unstable submanifolds $W_x^s, W_x^u \subseteq M$, of dimensions $\dim(X) - \text{ind}(x)$ and $\text{ind}(x)$, respectively. We can then formulate the Morse-Smale condition as in the gradient case:

Definition 2.9. *A Morse quasi-gradient vector field v is called Morse-Smale if for all stationary points x and y , the unstable manifold W_x^u and the stable manifold W_y^s intersect transversely.*

Morse-Smale quasi-gradients are a particular example of the *Morse-Smale vector fields*, which play an important role in the theory of dynamical systems. We recall their definition below. See [PdM82, Chapter 5] for an introduction to the subject.

Definition 2.10 ([PdM82], p.118-119). *Let X be a closed, smooth manifold. A smooth vector field v on X is called Morse-Smale if the following conditions are satisfied:*

- (a) *v has a finite number of critical elements (stationary points and closed orbits), all of which are hyperbolic.³*
- (b) *If σ_1 and σ_2 are critical elements of v , then the unstable manifold of σ_1 is transverse to the stable manifold of σ_2 ;*
- (c) *The set of nonwandering points for v is equal to the union of the critical elements of v . Here, $x \in X$ is called a wandering point for v if there exists a neighborhood U of x and a number $t_0 > 0$ such that $\Phi_t(U) \cap U = \emptyset$ for $|t| > t_0$, where $\{\Phi_t\}_{t \in \mathbb{R}}$ is the flow generated by v . Otherwise, we say that x is nonwandering.*

³See [PdM82, p.95] for the definition of hyperbolicity for closed orbits. Stable and unstable manifolds for closed orbits are defined in the obvious way; see [PdM82, p.98].

Morse-Smale vector fields satisfy structural stability, that is, a small perturbation of a Morse-Smale dynamical system is conjugate to the original system via a homeomorphism.

Lemma 2.11. *A Morse-Smale quasi-gradient vector field, in the sense of Definition 2.9, is a Morse-Smale vector field in the sense of Definition 2.10 and, furthermore, it has no closed orbits.*

Proof. Closed orbits do not exist because f decreases along flow lines. Parts (a) and (b) of Definition 2.10 are immediate consequences of Definition 2.9. Part (c), the fact that nonstationary points are wandering, follows since f decreases along flow lines. \square

Suppose v is a Morse-Smale quasi-gradient field. For every $x, y \in \mathfrak{C}$, the intersection

$$M(x, y) := W_x^u \cap W_y^s$$

is transverse, and hence a manifold of dimension $\text{ind}(x) - \text{ind}(y)$. We can construct a Morse complex (C_*, ∂) as in the gradient case. The groups C_* are generated by the stationary points, the grading is determined by the index, and the differential ∂ is given by the formula (1).

Proposition 2.12. *Let v be a Morse-Smale quasi-gradient vector field on a closed, oriented, smooth manifold X . Let (C_*, ∂) be the corresponding Morse complex. Then $\partial^2 = 0$, and the homology $H_*(C)$ is isomorphic to $H_*(M)$.*

Proof. In the gradient case, one way to prove that Morse homology recovers singular homology is by using the Conley index; cf. [Flo89, Sal90, Fra86, McC88] or [BH04, Chapter 7]. (See Section 2.4 below for the definition of the Conley index.) The basic idea is to find a Morse decomposition of the flow into attractor-repeller pairs. This produces a filtration of X by index pairs, such that the quotients are the Conley indices of the stationary points $x \in \mathfrak{C}$. These quotients are homotopy equivalent to spheres of dimensions $\text{ind}(x)$, and the connecting homomorphisms between their homology groups are given by the signed count of flow lines.

This proof can be adapted to the setting of Morse-Smale quasi-gradients. Indeed, one can construct Morse decompositions and connection matrices for general flows, as in the work of Franzosa [Fra86]. Further, Theorem 3.1 in [McC88] describes the neighborhood of a flow line between two hyperbolic points in a Morse-Smale (not necessarily gradient) flow. This shows that the connection matrices are given by counts of flow lines, i.e. that each gradient flow line contributes plus or minus 1 to the boundary operator. Orientations can be fixed as in [Flo89, Sal90]. \square

The rest of the discussion in Section 2.1 extends to quasi-gradients as well. The Morse-Smale condition can be phrased in terms of the surjectivity of the operator

$$L_\gamma(w) = \frac{Dw}{dt} + D(v)(w),$$

as in Lemma 2.2. Furthermore, the discussion from Section 2.2 carries over to this setting, so the moduli spaces $M(x, y)$ can be oriented using coherent orientations or specialized coherent orientations. The only caveat is that in the definition of coherent orientations, when we define the sets Σ_{γ^*TX} , we should allow the operators K^\pm to be hyperbolic, rather than conjugated self-adjoint.

2.4. Morse homology for isolated invariant sets. Suppose X is a (possibly non-compact) smooth, oriented manifold equipped with a Morse quasi-gradient vector field v . In general, there is no way of defining a Morse complex. Even if we assume the Morse-Smale condition on the unstable and stable manifolds, the resulting differential may not square to zero, because flow lines can escape to infinity.

Nevertheless, we have a Morse complex in a certain setting, described below. First, we recall the definitions of an isolated invariant set, index pair, and Conley index, following [Con78]. Fix a one-parameter subgroup $\{\Phi_t\}$ of diffeomorphisms of a manifold X . Given a compact subset $A \subseteq X$, define the compact subset

$$\text{Inv}(A, \Phi) = \{x \in A \mid \Phi_t(x) \in A \text{ for all } t \in \mathbb{R}\}.$$

Definition 2.13. A compact set $\mathcal{S} \subseteq X$ is called an isolated invariant set if there exists a compact $A \subseteq X$ with $\mathcal{S} = \text{Inv}(A, \Phi) \subseteq \text{int}(A)$. Such a set A is called an isolating neighborhood for \mathcal{S} .

Definition 2.14. An index pair (N, L) for \mathcal{S} is a pair of compact sets $L \subseteq N \subseteq X$ if

- a) $\text{Inv}(N - L, \Phi) = \mathcal{S} \subset \text{int}(N - L)$,
- b) for all $x \in N$, if there exists $t > 0$ such that $\Phi_t(x)$ is not in N , there exists $0 \leq \tau < t$ with $\Phi_\tau(x) \in L$ (L is an exit set for N),
- c) given $x \in L$, $t > 0$, if $\Phi_s(x) \subseteq N$ for all $0 \leq s \leq t$, then $\Phi_s(x)$ is in L for $0 \leq s \leq t$ (L is positively invariant in N).

Conley [Con78] showed that any isolated invariant set \mathcal{S} admits an index pair. The Conley index for an isolated invariant set \mathcal{S} , denoted $I(\Phi, \mathcal{S})$, is defined to be the pointed space $(N/L, [L])$. Its pointed homotopy type is an invariant of the triple (X, Φ_t, \mathcal{S}) . In fact, the Conley index is invariant under continuous deformations of the flow, as long as \mathcal{S} remains isolated in a suitable sense.

Going back to the case of a Morse-Smale quasi-gradient flow on X , suppose $\mathcal{S} \subseteq X$ is an isolated invariant set. Let $\mathfrak{C}[\mathcal{S}]$ be the set of stationary points of v that lie in \mathcal{S} . The set $\mathfrak{C}[\mathcal{S}]$ is finite, because \mathcal{S} is compact. For any $x, y \in \mathfrak{C}[\mathcal{S}]$, if a point of \mathcal{S} is on a flow line $\gamma \in \check{M}(x, y)$, then all the points on that flow line must be in \mathcal{S} . Further, the subsets

$$\check{M}[\mathcal{S}](x, y) = \{\gamma \in \check{M}(x, y) \mid \gamma \subset \mathcal{S}\}$$

are both open and closed in $\check{M}(x, y)$. We have

$$\mathcal{S} = \bigcup_{\substack{x, y \in \mathfrak{C}[\mathcal{S}] \\ \gamma \in \check{M}[\mathcal{S}](x, y)}} \gamma.$$

We can define a Morse complex $C[\mathcal{S}]$ to be freely generated by the points in $\mathfrak{C}[\mathcal{S}]$, with the differential given by counting only the flow lines in $\check{M}[\mathcal{S}](x, y)$. We still have $\partial^2 = 0$, and we obtain a Morse homology $H_*(C[\mathcal{S}])$.

The following theorem was proved by Floer in [Flo89] in the setting of Morse-Smale gradient flows. The extension to the quasi-gradient case goes along the same lines as the proof of Proposition 2.12. (In fact, the results of Franzosa and McCord used in that proof were originally phrased for more general flows.)

Theorem 2.15 (cf. [Flo89]). *Given an isolated invariant set \mathcal{S} in a Morse-Smale quasi-gradient flow ϕ on a manifold X , the Morse homology $H_*(C[\mathcal{S}])$ is isomorphic to the reduced homology of the Conley index of \mathcal{S} .*

Thus, if (N, L) is an index pair for \mathcal{S} and $L \subseteq N$ is a neighborhood deformation retract (which can easily be arranged), then the Morse homology is isomorphic to the relative homology $H_*(N, L)$.

2.5. Morse homology for manifolds with boundary. Let X be a smooth, oriented, compact manifold with boundary. We proceed to describe the construction of a Morse complex, $(\check{C}(X), \check{\partial})$, whose homology is isomorphic to $H_*(X; \mathbb{Z})$ (we will ignore the analogous constructions for $H_*(\partial M; \mathbb{Z})$ or $H_*(X, \partial X; \mathbb{Z})$). For gradient flows, this construction appeared in print in [KM07, Section 2.4], although it may have been known before. See also [Lau11], [BNR12] for related work. Our exposition follows [KM07] closely, except we use the more general quasi-gradient flows.

Fix a metric g and a Morse quasi-gradient vector field v on X such that g and v are respectively the restrictions of a metric and a vector field on the double of X , which are invariant under the obvious involution. Among the stationary points of v , there are those which occur in the interior of X ; their set is denoted \mathfrak{C}^o . The others are stationary points on ∂X . If $x \in \partial X$ is such a point, note that the normal vector to the boundary N_x is taken to its opposite under the involution, and the same holds for $dv(N_x)$. Hence, N_x is an eigenvector of dv , so it must live in either $T_x^s X$ or $T_x^u X$. In the first case we say that x is *boundary-stable*, and in the second that it is *boundary-unstable*. The set of boundary-stable stationary points is denoted \mathfrak{C}^s , and the set of boundary-unstable stationary points is denoted \mathfrak{C}^u .

We require that v satisfy the usual Morse-Smale condition, except in the so-called *boundary-obstructed* case, when x is boundary-stable and y is boundary-unstable. In the boundary-obstructed case, W_x^u and W_y^s are subsets of ∂X , so transversality in X is impossible. Therefore, we instead require that they intersect transversely in ∂X . When this happens, the dimension of $M(x, y)$ is in fact $\text{ind}(x) - \text{ind}(y) + 1$.

For any $x, y \in \mathfrak{C} = \mathfrak{C}^o \cup \mathfrak{C}^s \cup \mathfrak{C}^u$, we let $M^\partial(x, y) = M(x, y) \cap \partial X$ and $\check{M}^\partial(x, y) = M^\partial(x, y)/\mathbb{R}$. One can induce the same orientations on the moduli spaces as in the usual closed setting, except in the boundary-obstructed case, where we must also make use of an outward normal vector.

We let C^o , C^s , and C^u be the free Abelian groups generated by \mathfrak{C}^o , \mathfrak{C}^s , and \mathfrak{C}^u respectively. The chain groups we work with will be $\check{C}(X) = C^o \oplus C^s$. Although these chain groups do not incorporate the boundary-unstable stationary points, the differential does. The way that the flows are counted varies from the usual construction. We define two sets of maps

$$\partial_\varpi^\theta, \bar{\partial}_\varpi^\theta : C^\theta \rightarrow C^\varpi$$

for various pairs of $\theta, \varpi \in \{o, u, s\}$. If $(\theta, \varpi) \in \{(o, o), (o, s), (u, o), (u, s)\}$ and $x \in \mathfrak{C}^\theta$, define

$$\partial_\varpi^\theta x = \sum_{y \in \mathfrak{C}^\varpi} n(x, y)y,$$

where the summation is again over those y with $\dim \check{M}(x, y) = 0$.

If $\theta, \varpi \in \{s, u\}$ and $x \in \mathfrak{C}^\theta$, we instead define $\bar{\partial}_\varpi^\theta$ to count flows on ∂X :

$$\bar{\partial}_\varpi^\theta x = \sum_{y \in \mathfrak{C}^\varpi} \bar{n}(x, y)y,$$

where $\bar{n}(x, y)$ is the signed count of points (unparameterized flows) in zero-dimensional moduli spaces $\check{M}^\partial(x, y)$.

On \check{C} , we define $\check{\partial}$ by

$$(11) \quad \check{\partial} = \begin{bmatrix} \partial_o^o & -\partial_o^u \bar{\partial}_u^s \\ \partial_s^o & \bar{\partial}_s^s - \partial_s^u \bar{\partial}_u^s \end{bmatrix}.$$

Theorem 2.16 (Theorem 2.4.5 in [KM07]). *We still have $\check{\partial}^2 = 0$. Furthermore, $H_*(\check{C}, \check{\partial}) \cong H_*(X)$.*

Remark 2.17. In this paper, since $\bar{\partial}_s^u$ does not appear in $\check{\partial}$, we will not need the notation $M^{\partial}(x, y)$, since either $M^{\partial}(x, y)$ agrees with $M(x, y)$ or is empty, except when x is boundary-unstable and y is boundary-stable. In [KM07], the term $\bar{\partial}_s^u$ is needed for the proof that $\check{\partial}^2 = 0$. We keep the notation and discussion for consistency with [KM07].

2.6. Morse homology for manifolds with circle actions. A circle action is called *semifree* if it is free on the complement of the fixed points. For manifolds with semifree S^1 -actions, Kronheimer and Mrowka showed that the methods in the previous subsection can be applied to obtain Morse theoretic approximations to the S^1 -equivariant homology. We sketch their arguments here, following [KM07, Sections 2.5–2.6], but phrasing everything in terms of quasi-gradients rather than gradients.

Suppose that a closed Riemannian manifold X has a smooth, semifree S^1 -action by isometries, and let Q denote the fixed point set. Let $N(Q)$ be the normal bundle of Q , and $N^1(Q)$ the unit normal bundle. Observe that the S^1 -action gives $N(Q)$ the structure of a complex vector bundle. In order to make X/S^1 into a manifold, we resolve the singularity at Q by considering the blow-up

$$X^\sigma = (X - Q) \cup (N^1(Q) \times [0, \epsilon)),$$

where we have identified $N^1(Q) \times (0, \epsilon)$ with $N(Q) - Q \subset X$. Note that X^σ/S^1 is a smooth manifold with boundary $N^1(Q)/S^1$.

A smooth, S^1 -invariant vector field \tilde{v} on X induces a vector field v on $(X - Q)/S^1$. In turn, this extends naturally to a smooth vector field v^σ on X^σ/S^1 ; cf. [KM07, Lemma 2.5.2]. Note that \tilde{v} must be tangent to Q , and hence v^σ is tangent to the boundary.

Precisely, the dynamics of v^σ on the boundary $N^1(Q)/S^1 = \partial(X^\sigma/S^1)$ can be described as follows. A point on the boundary can be written as $(q, [\phi])$ with $q \in Q$ and $\phi \in N_q^1(Q)$. As discussed, $N_q(Q)$ has a complex structure; we let $\langle \phi \rangle^\perp$ be the complex orthogonal complement to ϕ in $N_q(Q)$. Decompose the tangent space to X^σ/S^1 at q as

$$(12) \quad T_q Q \oplus \langle \phi \rangle^\perp \oplus \mathbb{R},$$

where \mathbb{R} is the direction normal to the boundary.

The covariant derivative $(\nabla \tilde{v})_q : T_q X \rightarrow T_q X$ is S^1 -equivariant, and hence takes the normal direction $N_q Q \subset T_q X$ to $N_q Q$. Let

$$(13) \quad L_q := (\nabla \tilde{v})|_{N_q Q}.$$

With respect to the decomposition (12), let us write

$$v^\sigma(q, [\phi]) = (\tilde{v}(q), \mathbb{L}_q \phi, 0).$$

Here, when $\tilde{v}(q) = 0$, the second term equals

$$\mathbb{L}_q \phi = L_q \phi - \operatorname{Re} \langle \phi, L_q \phi \rangle \phi.$$

Thus, the stationary points of v^σ on the boundary are the pairs $(q, [\phi])$, where q is a zero of $\tilde{v}|_Q$ and ϕ is an eigenvector of L_q .

Furthermore, the flow associated to v^σ on the boundary is given by the equations:

$$(14) \quad \frac{dq}{dt} + \tilde{v}(q(t)) = 0,$$

$$(15) \quad q^*(\nabla) + (L_{q(t)} - \operatorname{Re}\langle \phi(t), L_{q(t)}\phi(t) \rangle \phi(t))dt = 0,$$

where $|\phi(t)| = 1$ for all t .

Definition 2.18. A smooth, S^1 -invariant vector field \tilde{v} on X is called a Morse equivariant quasi-gradient if the following conditions are satisfied:

- (a) All stationary points of v on $(X - Q)/S^1$ are hyperbolic.
- (b) All stationary points of $\tilde{v}|_Q$ are hyperbolic.
- (c) At each stationary point q of $\tilde{v}|_Q$, the operator $L_q : N_q Q \rightarrow N_q Q$ is self-adjoint, and admits a basis of eigenvectors $\phi_1(q), \phi_2(q), \dots, \phi_n(q)$ with corresponding eigenvalues $\lambda_1(q), \lambda_2(q), \dots, \lambda_n(q)$ such that

$$\lambda_1(q) < \lambda_2(q) < \dots < \lambda_n(q)$$

and $\lambda_i(q) \neq 0$ for any i .

- (d) There exists a smooth, S^1 -equivariant function $\tilde{f} : X \rightarrow \mathbb{R}$ such that $d\tilde{f}(\tilde{v}) \geq 0$ at all $x \in X$, with equality holding if and only if $\tilde{v}(x) = 0$.

Lemma 2.19. Parts (a), (b) and (c) in Definition 2.18, taken together, are equivalent to asking for all the stationary points of v^σ on X^σ/S^1 to be hyperbolic, and for those on the boundary to give rise to operators L_q that are self-adjoint.

Furthermore, if \tilde{v} is as in Definition 2.18, and $(q, [\phi_i(q)])$ is a stationary point of v^σ on the boundary, then its index is given by

$$\operatorname{ind}(q, [\phi_i(q)]) = \begin{cases} \operatorname{ind}_Q(q) + 2i - 2 & \text{if } \lambda_i(q) > 0, \\ \operatorname{ind}_Q(q) + 2i - 1 & \text{if } \lambda_i(q) < 0. \end{cases}$$

Proof. Away from the boundary, a stationary point of v^σ is just a stationary point of v , and the two notions of hyperbolicity (and index) correspond. With regard to stationary points on the boundary and their indices, see [KM07, Proof of Lemma 2.5.5] for the case of gradient fields. The arguments there extend to the quasi-gradient setting without difficulty. \square

Remark 2.20. In part (c) of Definition 2.18, the condition that L_q is self-adjoint (and has real eigenvalues) is not strictly necessary. We could have only asked for L_q to have a complex basis of eigenvectors $\phi_1(q), \phi_2(q), \dots, \phi_n(q)$ with corresponding eigenvalues $\lambda_1(q), \lambda_2(q), \dots, \lambda_n(q)$ such that

$$\operatorname{Re} \lambda_1(q) < \operatorname{Re} \lambda_2(q) < \dots < \operatorname{Re} \lambda_n(q)$$

and $\operatorname{Re} \lambda_i(q) \neq 0$ for any i . With this weaker requirement, the results of this section would still hold, but the proof of Lemma 2.23 below would be more complicated. In the case of interest to us in this paper, the L_q are self-adjoint, so we decided to include this condition in the definition.

Lemma 2.21. Let \tilde{v} be a Morse equivariant quasi-gradient vector field, and let f be a function as in part (d) of Definition 2.18. Then, any stationary point of \tilde{v} is a critical point of f .

Proof. This is similar to the proof of Lemma 2.7. \square

Given $x \in X$, we write $[x]$ for its S^1 -equivalence class, i.e. its projection to the singular space X/S^1 .

Lemma 2.22. *Let \tilde{v} be a Morse equivariant quasi-gradient vector field on X , as in Definition 2.18. Let $\gamma : \mathbb{R} \rightarrow X$ be a flow line of \tilde{v} . Then, $\lim_{t \rightarrow -\infty} [\gamma(t)]$ and $\lim_{t \rightarrow +\infty} [\gamma(t)]$ exist in X/S^1 , and they are both projections of stationary points of \tilde{v} .*

Proof. This is similar to the proof of Lemma 2.8. \square

Suppose \tilde{v} is a Morse equivariant quasi-gradient vector field on X , with \tilde{f} as in part (d) of Definition 2.18. Observe that \tilde{f} induces a function $f : X/S^1 \rightarrow \mathbb{R}$ which is smooth except on Q/S^1 . Using f we see that v^σ is a Morse quasi-gradient vector field on the interior of X^σ/S^1 .

If we let $f^\sigma : X^\sigma/S^1 \rightarrow \mathbb{R}$ be the composition of the blow-down map with f , then f^σ is smooth, and we have

$$(16) \quad df^\sigma(v^\sigma) \geq 0.$$

However, \tilde{f} is constant on the fibers $N_q^1(Q)/S^1$. Hence, equality happens in (16) not only at the stationary points of v^σ , but also at all $(q, [\phi]) \in \partial(X^\sigma/S^1)$ such that $\tilde{v}(q) = 0$ (and ϕ does not have to be an eigenvector). Therefore, v^σ is not naturally a Morse quasi-gradient on all of X^σ/S^1 ; or, at least, we cannot use f^σ to argue that it is.

Nevertheless, we will still be able to do Morse homology using the flow of v^σ . To start with, observe that there are no closed orbits in this flow: In view of (16), the only such orbits would have to be contained in a fiber $N_q^1(Q)/S^1$, where $\tilde{v}(q) = 0$. From (15) we see that the flow on the projective space $N_q^1(Q)/S^1$ is the projection of a linear flow on $N_q Q$, which is in fact a gradient flow. This shows that there are no closed orbits in that fiber.

Furthermore, we have the analogue of Lemmas 2.8 and 2.22:

Lemma 2.23. *Let \tilde{v} be a Morse equivariant quasi-gradient vector field on X , as in Definition 2.18. Let $\gamma : \mathbb{R} \rightarrow X^\sigma/S^1$ be a flow line of v^σ . Then, $\lim_{t \rightarrow -\infty} \gamma(t)$ and $\lim_{t \rightarrow +\infty} \gamma(t)$ exist in X^σ/S^1 , and they are both stationary points of v^σ .*

Proof. By Lemma 2.22, we already know that the projection of γ to the blow-down X/S^1 limits to two stationary points. The projection is one-to-one away from the boundary, so it suffices to study the case of a trajectory that (in the blow-down) limits to a stationary point q_0 in the fixed point set Q . The limit can be at either $-\infty$ and $+\infty$. Without loss of generality, we consider the case of $+\infty$, and focus on the half-trajectory

$$\gamma_+ : [0, \infty) \rightarrow X^\sigma/S^1, \quad \gamma_+(t) = \gamma(t).$$

The idea is that, in the blow-up, the flow of v^σ near the boundary $N^1(Q)/S^1$ is approximated by the flow on the boundary, which is given by (14)-(15), and whose behavior we understand.

A suitably small neighborhood $V \subset X^\sigma/S^1$ of the boundary $\partial(X^\sigma/S^1) = N^1(Q)/S^1$ can be identified with the normal bundle to the boundary. Hence, a point $v \in V$ can be written as a triple $(q, s, [\phi])$, with $q \in Q$, $s \geq 0$, and $\phi \in N_q^1(Q)$, normalized so that $|\phi| = 1$. (Here, s represents the distance to the boundary.) With L_q as in (13), we introduce the functional

$$\Lambda : V \rightarrow \mathbb{R}, \quad \Lambda(q, s, [\phi]) = \langle \phi, L_q \phi \rangle.$$

The stationary points of v^σ on the boundary correspond to $s = 0$ and $d\Lambda = 0$.

Consider the fiber $F = N_{q_0}^1(Q)/S^1$ over q_0 . The restriction of v^σ to F is the gradient of $\frac{1}{2}\Lambda|_F$. By part (c) of Definition 2.18, $v^\sigma|_F$ has n stationary points, corresponding to the eigenvalues $\lambda_i = \lambda_i(q_0)$ of L_q , such that

$$\lambda_1 < \lambda_2 < \cdots < \lambda_n.$$

From our assumption, we know that given any neighborhood U of F , there exists t_0 such that $\gamma(t) \in U$ for $t > t_0$. In particular, all the accumulation points of the half-trajectory γ_+ must be contained in F . We seek to show that γ_+ has a unique accumulation point, and that point is a stationary point of v^σ .

Note that if $x \in F$ is an accumulation point of γ_+ , then all the points on the flow trajectory ζ through x are also accumulation points. Indeed, if $\{\Phi_t\}$ denotes the flow of v^σ , and $\gamma(t_n) \rightarrow x = \zeta(0)$, then $\gamma(t_n + t) = \Phi_t(\gamma(t_n)) \rightarrow \Phi_t(x) = \zeta(t)$. The trajectory ζ is contained in F , where the flow is a gradient flow, and therefore ζ limits to two stationary points $x = (q_0, 0, [\phi_i])$ and $y = (q_0, 0, [\phi_j])$ at $\mp\infty$. Moreover, x and y have to be accumulation points for γ_+ as well, and we have $x = y$ if and only if ζ is stationary.

Therefore, we are left to show that γ_+ cannot have two different stationary points $x, y \in F$ as accumulation points. Suppose that were the case, and let $\lambda_i > \lambda_j$ be the eigenvalues corresponding to x and y . These are also the values of Λ at x and y . Pick an intermediate value $\lambda \in (\lambda_j, \lambda_i)$, such that $\lambda \neq \lambda_k$ for any k . Since v^σ restricts to be the gradient of $\frac{1}{2}\Lambda$ on F , we have

$$d\Lambda(v^\sigma) > 0 \text{ on } \Lambda^{-1}(\lambda) \cap F.$$

Since $\Lambda^{-1}(\lambda) \cap F$ is compact, we can find a neighborhood $U \subset V$ of F such that

$$(17) \quad d\Lambda(v^\sigma) > 0 \text{ on } \Lambda^{-1}(\lambda) \cap U.$$

Because x and y are accumulation points of γ_+ , it must be that the half-trajectory γ_+ intersects the intermediate level set $\Lambda^{-1}(\lambda)$ infinitely many times. Furthermore, we know that after a certain time t_0 , the trajectory γ_+ is contained in U . However, this contradicts (17), which says that Λ decreases every time the trajectory goes through $\Lambda^{-1}(\lambda) \cap U$. The conclusion follows. \square

Since the stationary points of v^σ are hyperbolic, they admit stable and unstable manifolds; cf. the discussion in Section 2.3. Further, we can separate the stationary points on the boundary into stable and unstable, and then define the notion of boundary-obstructed trajectories, as in Section 2.5.

Definition 2.24. *A Morse equivariant quasi-gradient vector field \tilde{v} on X is called Morse-Smale if the induced vector field v^σ on X^σ/S^1 satisfies the Morse-Smale condition for boundary-unobstructed trajectories; and the Morse-Smale condition inside $\partial(X^\sigma/S^1)$ for the boundary-obstructed trajectories.*

If we have a Morse-Smale equivariant quasi-gradient vector field, the constructions in Section 2.5 carry through using v^σ (even though v^σ is not a quasi-gradient vector field in a natural way). In particular, Lemma 2.23 implies that v^σ is a Morse-Smale vector field as in Definition 2.10; compare Lemma 2.11. Thus, we obtain a Morse complex

$$(\check{C}(X^\sigma/S^1), \check{\partial})$$

which computes the homology of X^σ/S^1 .

The space X^σ/S^1 can be viewed as an approximation to the homotopy quotient $X//S^1 := X \times_{S^1} ES^1$. Precisely, let n be the connectivity of the pair $(X, X - Q)$, that is, the largest

j such that $(X, X - Q)$ is j -connected. Note that if the real codimension of Q in X is $2c$, then $n \geq 2c - 1$. Since X^σ is S^1 -equivariantly homotopy equivalent to $X - Q$, we have an isomorphism in homology

$$(18) \quad H_j(X^\sigma/S^1) \cong H_j(X//S^1)$$

for all $j \leq n - 1$. The homology $H_*(X//S^1) \cong H_*^{S^1}(X)$ is the S^1 -equivariant Borel homology of X .

2.7. The U -action in Morse homology. We keep the same setting as in the previous subsection. The equivariant homology $H_*^{S^1}(X)$ admits a natural $\mathbb{Z}[U]$ -module structure, given by cap products with the elements of $H_{S^1}^*(pt) \cong \mathbb{Z}[U]$. (The action of U decreases degree by two.) Our goal here is to explain how this module structure can be approximated in terms of Morse theory. The discussion is modeled on the infinite-dimensional case presented in [KMOS07, Section 4.11].

The circle action on X^σ produces a natural complex line bundle E^σ on X^σ/S^1 . Let ζ be a generic, smooth section of E^σ , such that ζ is transverse to the 0-section and further, ζ restricted to the boundary, $N^1(Q)/S^1$ is transverse to restriction of the 0-section to the boundary. Note that ζ inherits a canonical orientation from X^σ/S^1 . Let \mathcal{Z} denote the zero set of ζ , which is a manifold with boundary $\partial\mathcal{Z} = \mathcal{Z} \cap (N^1(Q)/S^1)$. From the orientation of ζ and X^σ/S^1 , we obtain an orientation on \mathcal{Z} .

Notice that a trajectory γ of v^σ is determined by its value at time $t = 0$. Thus, we can define cut-down moduli spaces by intersecting $M(x, y)$ with \mathcal{Z} at time $t = 0$:

$$M(x, y) \cap \mathcal{Z} := \{\gamma \in M(x, y) \mid \gamma(0) \in \mathcal{Z}\}.$$

Given $(\theta, \varpi) \in \{(o, o), (o, s), (u, o), (u, s)\}$, we define maps of the form $m_{\varpi}^\theta : \check{C}_*^\theta(X^\sigma/S^1) \rightarrow \check{C}_{*-2}^\varpi(X^\sigma/S^1)$ by counting flows that intersect \mathcal{Z} :

$$m_{\varpi}^\theta(x) = \sum_{y \in \mathfrak{c}^\varpi} \#(M(x, y) \cap \mathcal{Z}) \cdot y.$$

Note that the conditions on the gradings of x and y guarantee that the cut-down moduli spaces being counted are 0-dimensional.

For $\theta, \varpi \in \{u, s\}$, we also define analogous maps that count intersections in the boundary of X :

$$\bar{m}_{\varpi}^\theta x = \sum_{y \in \mathfrak{c}^\varpi} \#(M(x, y) \cap \partial\mathcal{Z}) \cdot y,$$

where we only consider y where the relevant cut-down moduli space is 0-dimensional. Further, recall the shift in gradings by one in the boundary-obstructed case. We define a chain map $\check{m} : \check{C}_*(X^\sigma/S^1) \rightarrow \check{C}_{*-2}(X^\sigma/S^1)$ by

$$(19) \quad \check{m} = \begin{bmatrix} m_o^o & -m_o^u \bar{\partial}_u^s - \partial_o^u \bar{m}_u^s \\ m_s^o & \bar{m}_s^s - m_s^u \bar{\partial}_u^s - \partial_s^u \bar{m}_u^s \end{bmatrix}.$$

The map induced by \check{m} on the homology $H_*(X^\sigma/S^1)$ is exactly the cap product with $c_1(E^\sigma)$. This recovers the $\mathbb{Z}[U]$ -module structure on $H_*(X^\sigma/S^1) = H_*^{S^1}(X^\sigma)$. Note that the isomorphism (18) discussed in the previous subsection,

$$(20) \quad H_{\leq n-1}(X^\sigma/S^1) \cong H_{\leq n-1}(X//S^1),$$

is actually an isomorphism of $\mathbb{Z}[U]$ -modules. Thus, Morse theory tells us the action of $\mathbb{Z}[U]$ on $H_*^{S^1}(X)$ in degrees up to $n - 1$.

2.8. Combined generalizations. The versions of Morse homology described in Sections 2.4 and 2.5 can be combined as follows. Let X be a (possibly non-compact) manifold with boundary. Fix a metric g and a Morse quasi-gradient vector field v on X as in Section 2.5, and let $\mathcal{S} \subseteq X$ be an isolated invariant set of the resulting flow. Although our definitions in Section 2.4 were for flows on manifolds without boundary, Conley index theory easily extends to the boundary case, provided that the gradient vector field is tangent to the boundary; compare [Hel09, Section 3.1.2]. The only caveat is that in Definitions 2.13 and 2.14, the interior of a subset $A \subseteq X$ should be defined as in point-set topology, without regard to the structure of X as a manifold-with-boundary. With this in mind, the Conley index of \mathcal{S} is defined as before, to be the quotient N/L of an index pair (N, L) for \mathcal{S} . Alternatively, we could consider the double $D(X)$ of X with its $\mathbb{Z}/2$ -action, and appeal to the equivariant Conley index theory developed in [Flo87, Pru99]. This guarantees the existence of a $\mathbb{Z}/2$ -equivariant Conley index of $D(\mathcal{S})$ on $D(X)$, well-defined up to $\mathbb{Z}/2$ -equivariant homotopy equivalence. In particular, its quotient by $\mathbb{Z}/2$, which we take to be the Conley index of \mathcal{S} on X , is well-defined up to homotopy equivalence. This is equivalent to [Hel09, Definition 3.1.32].

We now impose a Morse-Smale condition for the trajectories in \mathcal{S} as in Section 2.5, where in the boundary-obstructed case we only require transversality inside ∂X . We then construct a complex $\check{C}[\mathcal{S}]$ using the same formula (11), but involving only the stationary points and the flow trajectories in \mathcal{S} . The homology of $\check{C}[\mathcal{S}]$ will be the reduced homology of the Conley index associated to \mathcal{S} .

Starting from this, we can also combine the construction in Section 2.4 with those in Sections 2.6-2.7. Suppose that X is a (possibly non-compact) Riemannian manifold X with a smooth, semifree S^1 -action by isometries. (In our applications, X will be a vector space of the form $\mathbb{R}^m \oplus \mathbb{C}^n$, with the linear S^1 -action.) Let \tilde{v} be a smooth vector field on X , and $\mathcal{S} \subseteq X$ be an S^1 -invariant, isolated invariant set in the flow of \tilde{v} . Let $Q, X^\sigma, v^\sigma, L_q$ be constructed as in Section 2.6. Given a closed S^1 -invariant subset $M \subset X$, we will denote by M^σ the closure of the preimage of $M - Q$ in X^σ . We seek to form a complex $\check{C}(X^\sigma/S^1)[\mathcal{S}]$ using only trajectories in \mathcal{S}^σ/S^1 . In order to do this, we require that we can find an S^1 -invariant isolating neighborhood A of \mathcal{S} such that the restriction of \tilde{v} to A is a Morse-Smale equivariant quasi-gradient vector field in the sense of Definition 2.24.

Under these assumptions, we obtain the desired Morse complex $\check{C}(X^\sigma/S^1)[\mathcal{S}]$ using trajectories in A^σ/S^1 . Observe that \mathcal{S}^σ/S^1 is an isolated invariant set for the flow of v^σ on X^σ/S^1 ; let $I(\mathcal{S}^\sigma/S^1)$ be its Conley index. We have:

$$(21) \quad H_*(\check{C}(X^\sigma/S^1)[\mathcal{S}]) \cong \tilde{H}_*(I(\mathcal{S}^\sigma/S^1)).$$

On the other hand, \mathcal{S} is an isolated invariant set itself, and is fixed by the S^1 -action. We may choose an index pair which is S^1 -invariant [Flo87, Pru99], and we thus get an S^1 -equivariant Conley index $I_{S^1}(\mathcal{S})$. Let $(I_{S^1}(\mathcal{S}))^{S^1}$ be its fixed point set. The Morse homology in (21) approximates the reduced equivariant homology of $I_{S^1}(\mathcal{S})$, in the sense that:

$$(22) \quad H_{\leq n-1}(\check{C}(X^\sigma/S^1)[\mathcal{S}]) \cong \tilde{H}_{\leq n-1}^{S^1}(I_{S^1}(\mathcal{S})),$$

where n is the connectivity of the pair $(I_{S^1}(\mathcal{S}), (I_{S^1}(\mathcal{S}) - (I_{S^1}(\mathcal{S}))^{S^1}) \cup *)$, with $*$ denoting the basepoint in $I_{S^1}(\mathcal{S})$; compare (20). Moreover, after choosing a suitable section ζ of E^σ as in Section 2.7, the isomorphism in (22) becomes one of $\mathbb{Z}[U]$ -modules.

Finally, recall that in Sections 2.1 and 2.2 we described alternative ways of stating the Morse-Smale condition and of constructing orientations. These descriptions apply equally well to the more general settings discussed here.

3. THE SEIBERG-WITTEN FLOER SPECTRUM

We review here the construction of the Seiberg-Witten Floer spectrum $\text{SWF}(Y, \mathfrak{s})$, following [Man03].

3.1. The configuration space and the gauge group action. We will be studying the Seiberg-Witten equations on a tuple $(Y, g, \mathfrak{s}, \mathbb{S})$, where Y is a rational homology three-sphere, g is a metric on Y , \mathfrak{s} is a Spin^c structure on Y , and \mathbb{S} is a spinor bundle for \mathfrak{s} . We choose a flat Spin^c connection A_0 on \mathbb{S} which gives an affine identification of $\Omega^1(Y; i\mathbb{R})$ with Spin^c connections on \mathbb{S} .

We will be doing analysis on the configuration space

$$\mathcal{C}(Y) = \Omega^1(Y; i\mathbb{R}) \oplus \Gamma(\mathbb{S}).$$

Of course, $\mathcal{C}(Y)$ also depends on \mathfrak{s} , but we omit it from the notation. The gauge group $\mathcal{G} = \mathcal{G}(Y) := C^\infty(Y, S^1)$ acts on $\mathcal{C}(Y)$ by $u \cdot (a, \phi) = (a - u^{-1}du, u \cdot \phi)$. Since $b_1(Y) = 0$, each $u \in \mathcal{G}$ can be written as e^f for some $f : Y \rightarrow i\mathbb{R}$. We define the *normalized gauge group* \mathcal{G}^o to consist of those $u = e^f \in \mathcal{G}$ such that $\int_Y f = 0$.

For any integer k , following [KM07], we let H_k denote the completion of a subspace $H \subseteq \mathcal{C}(Y)$ with respect to the L_k^2 Sobolev norm. In particular, the completion of W is denoted W_k . The Sobolev norm is defined in the standard way using the L^2 norms of iterated gradients ∇^j , as in [KM07, Section 5.1]. For $j \leq k$, define \mathcal{T}_j as the L_j^2 completion of $T\mathcal{C}_k(Y)$. (In particular, \mathcal{T}_k is $T\mathcal{C}_k(Y)$.)

3.2. Coulomb slices. We have a *global Coulomb slice*⁴:

$$W = \ker d^* \oplus \Gamma(\mathbb{S}) \subset \mathcal{C}(Y),$$

where d^* is meant to act on imaginary 1-forms. Given $(a, \phi) \in \mathcal{C}(Y)$, there is a unique element of W which is obtained from (a, ϕ) by a normalized gauge transformation; this element is called the *global Coulomb projection* of (a, ϕ) . Explicitly, the global Coulomb projection of (a, ϕ) is

$$(23) \quad \Pi^{\text{gC}}(a, \phi) = (a - df, e^f \phi),$$

where $f : Y \rightarrow i\mathbb{R}$ is such that $d^*(a - df) = 0$ and $\int_Y f = 0$; that is, $f = Gd^*a$, where G is the Green's operator of $\Delta = d^*d$. For future reference, let us denote by

$$\pi : \Omega^1(Y; i\mathbb{R}) \rightarrow \ker d^*$$

the L^2 -orthogonal projection, given by $\pi(a) = a - df = a - dGd^*a$.

The derivative of Π^{gC} is called the *infinitesimal global Coulomb projection*, given by

$$(\Pi_*^{\text{gC}})_{(a, \phi)}(b, \psi) = (b - dGd^*b, e^{Gd^*a}(\psi + (Gd^*b)\phi)).$$

In particular, if (a, ϕ) happens to be already in W , we have

$$(24) \quad (\Pi_*^{\text{gC}})_{(a, \phi)}(b, \psi) = (b - d\xi, \psi + \xi\phi) = (\pi(b), \psi + \xi\phi) \in T_{(a, \phi)}W,$$

where $\xi = Gd^*b$.

Analogous to the definition of \mathcal{T}_j , let $\mathcal{T}_j^{\text{gC}}$ be the L_j^2 completion of the tangent bundle to W_k , namely the trivial vector bundle with fiber W_j over W_k . We keep the notation $\mathcal{T}_j^{\text{gC}}$

⁴The global Coulomb slice was denoted V in [Man03]. We switched to W in order to avoid confusion with the spaces denoted $\mathcal{V}(Z)$ in [KM07], which will also appear in this paper.

(rather than just $W_k \times W_j$) to emphasize the bundle structure; this will be convenient when we discuss bundle decompositions.

Let us mention the following lemma, which will be of use to us later:

Lemma 3.1. *Let $k \geq 2$. View the infinitesimal global Coulomb projection as a section of the bundle $\text{Hom}(TC(Y), TW)$ over W , i.e., a map from W to $\text{Hom}(\mathcal{C}(Y), W)$, given by*

$$(a, \phi) \mapsto ((b, \psi) \mapsto (\Pi_*^{\text{gC}})_{(a, \phi)}(b, \psi)).$$

Then, this map extends to smooth maps between the Sobolev completions

$$W_k \rightarrow \text{Hom}(\mathcal{C}_j(Y), W_j)$$

for all $-k \leq j \leq k$.

Proof. Recall the formula (24):

$$(\Pi_*^{\text{gC}})_{(a, \phi)}(b, \psi) = (b - dGd^*b, \psi + (Gd^*b)\phi).$$

Observe that d and d^* decrease Sobolev coefficients by one, the Green operator G increases them by 2. Because Sobolev multiplication $L_k^2 \times L_j^2 \rightarrow L_j^2$ induces a smooth map $L_k^2 \rightarrow \text{Hom}(L_j^2, L_j^2)$, the desired map is smooth. \square

Infinitesimally, we can consider a different slice to the gauge action, the one perpendicular to the orbits in the L^2 metric. This is the *local Coulomb slice* at $(a, \phi) \in \mathcal{C}(Y)$, denoted $\mathcal{K}_{(a, \phi)}$, and consisting of those tangent vectors $(b, \psi) \in T_{(a, \phi)}\mathcal{C}(Y)$ such that

$$(25) \quad -d^*b + i \text{Re}\langle i\phi, \psi \rangle = 0.$$

Away from the reducibles, we see that \mathcal{T}_k splits into a direct sum of two bundles:

$$\mathcal{T}_k = \mathcal{J}_k \oplus \mathcal{K}_k,$$

where \mathcal{J}_k consists of the vectors tangent to the \mathcal{G}_{k+1} orbits, and \mathcal{K}_k is the completion of the local Coulomb slice. Note that the local Coulomb slice does not form a bundle over the entire configuration space, since the local Coulomb slice is “bigger” at reducibles.

Given any $(b, \psi) \in T_{(a, \phi)}\mathcal{C}(Y)$, we define the (infinitesimal) *local Coulomb projection* of (b, ψ) to be

$$\Pi_{(a, \phi)}^{\text{lc}}(b, \psi) := (b - d\zeta, \psi + \zeta\phi),$$

where $\zeta : Y \rightarrow i\mathbb{R}$ is, for $\phi \neq 0$, the unique function such that

$$(26) \quad -d^*(b - d\zeta) + i \text{Re}\langle i\phi, \psi + \zeta\phi \rangle = 0.$$

The existence and uniqueness of such a ζ follow from [KM07, Proposition 9.3.4]. When $\phi = 0$, we again ask for (31) to be satisfied, but to guarantee uniqueness we also impose the condition $\int_Y \zeta = 0$.

Note that both types of Coulomb slices are also mentioned by Kronheimer and Mrowka; see [KM07, Sections 9.3 and 9.6].

For our purposes, we will also need the *enlarged local Coulomb slice*, which consists of vectors that are only required to be perpendicular to the orbits of the normalized gauge group action. We denote this by $\mathcal{K}_{(a, \phi)}^e$. A vector $(b, \psi) \in T_{(a, \phi)}\mathcal{C}(Y)$ is in $\mathcal{K}_{(a, \phi)}^e$ if and only if $-d^*b + i \text{Re}\langle i\phi, \psi \rangle$ is a constant function. Equivalently, we can write this condition as

$$(27) \quad -d^*b + i \text{Re}\langle i\phi, \psi \rangle^\circ = 0.$$

Here, and later in the paper, given a smooth function $f : Y \rightarrow \mathbb{C}$, we denote by $\mu_Y(f)$ the average value of f over the 3-manifold Y :

$$(28) \quad \mu_Y(f) = \frac{1}{\text{vol}(Y)} \int_Y f$$

and set

$$(29) \quad f^\circ = f - \mu_Y(f),$$

so that $\int_Y f^\circ = 0$.

We remark that the reason why we did not add the superscript \circ to d^*b in (27) is because $\int_Y d^*b = 0$, so $(d^*b)^\circ = d^*b$. Note also that unlike the local Coulomb slice, the enlarged local Coulomb slices do produce a bundle over the configuration space.

Given any $(b, \psi) \in T_{(a, \phi)}\mathcal{C}(Y)$, we define the *enlarged local Coulomb projection* of (b, ψ) to be

$$(30) \quad \Pi_{(a, \phi)}^{\text{elC}}(b, \psi) := (b - d\zeta, \psi + \zeta\phi),$$

where $\zeta : Y \rightarrow i\mathbb{R}$ is such that $\int_Y \zeta = 0$ and

$$(31) \quad -d^*(b - d\zeta) + i \text{Re}\langle i\phi, \psi + \zeta\phi \rangle^\circ = 0.$$

The fact that Π^{elC} is well-defined is established by the following lemma.

Lemma 3.2. (a) Fix $k, j \in \mathbb{Z}$ with $k \geq 2$ and $-k < j \leq k$. Then, for any $x = (a, \phi) \in \mathcal{C}_k(Y)$ and $(b, \psi) \in \mathcal{T}_{j, x}$, there is a unique $\zeta \in L_{j+1}^2(Y; i\mathbb{R})$ satisfying $\int_Y \zeta = 0$ and (31). Further, if $d^*b = 0$, then $\zeta \in L_{j+2}^2(Y; i\mathbb{R})$.

(b) If $d^*b = 0$ and the L_k^2 norm of ϕ is bounded above by a constant R , $k \geq 3$ and $|j| \leq k - 1$, then there exists a constant $C(R) > 0$ such that

$$\|\zeta\|_{L_{j+2}^2} \leq C(R) \cdot \|\psi\|_{L_j^2}.$$

Proof. Consider the direct sum decomposition

$$(32) \quad C^\infty(Y; i\mathbb{R}) = (\text{im } d^*) \oplus \mathbb{R}$$

into functions that integrate to zero and constant functions. This induces a similar decomposition on the L_j^2 Sobolev completions.

Let $\Delta = d^*d$ denote the (geometer's) Laplacian on imaginary-valued functions. Consider the linear operator between Sobolev completions

$$E_\phi : L_{j+1}^2(Y; i\mathbb{R}) \rightarrow L_{j-1}^2(Y; i\mathbb{R})$$

given by

$$E_\phi(\zeta) = (\Delta\zeta + (|\phi|^2\zeta)^\circ) + \int_Y \zeta.$$

With respect to the decomposition (32) for L_{j-1}^2 , note that the expression $\Delta\zeta + (|\phi|^2\zeta)^\circ$ lands in the first summand, and $\int_Y \zeta$ in the second summand.

Equation (31) together with the condition $\int_Y \zeta = 0$ can be written as

$$E_\phi(\zeta) = (d^*b - i \text{Re}\langle i\phi, \psi \rangle + i\mu_Y(\text{Re}\langle i\phi, \psi \rangle)) + 0.$$

Observe that the right hand side lives in $L_{j-1}^2(Y; i\mathbb{R})$ in general, and in $L_j^2(Y; i\mathbb{R})$ when $d^*b = 0$.

We need to show that E_ϕ is invertible. Observe that E_ϕ is a compact deformation of the operator $E_0 = \Delta + \int_Y$. The latter is invertible, and in particular Fredholm of index zero.

Hence, E_ϕ is also Fredholm of index zero. To show that it is invertible, it suffices to show that it has no kernel. Indeed, suppose $\zeta \in \ker(E_\phi)$. Then, $\int_Y \zeta = 0$ and

$$(33) \quad 0 = \int_Y \langle \Delta \zeta + |\phi|^2 \zeta - \mu_Y(|\phi|^2 \zeta), \zeta \rangle$$

$$(34) \quad = \int_Y |d\zeta|^2 + \int_Y |\phi|^2 |\zeta|^2 - 0.$$

This implies $d\zeta = 0$, and since $\int_Y \zeta = 0$, we get $\zeta = 0$. We conclude that E_ϕ is injective and hence invertible. This proves part (a).

For part (b), note that if $k \geq 3$, the map

$$L_{k-1}^2(Y; \mathbb{S}) \rightarrow \text{Hom}(L_{j+2}^2(Y; i\mathbb{R}), L_j^2(Y; i\mathbb{R})), \quad \phi \mapsto E_\phi$$

is continuous (and lands in invertible operators) for $|j| \leq k-1$. Hence, the map $\phi \mapsto E_\phi^{-1}$ is also continuous. Since the ball of radius R in L_k^2 is precompact in L_{k-1}^2 , the resulting collection of E_ϕ^{-1} is also precompact, and hence bounded, in $\text{Hom}(L_{j+2}^2(Y; i\mathbb{R}), L_j^2(Y; i\mathbb{R}))$. This gives the desired bounds. \square

Observe that, for $(a, \phi) \in W$, we can view the (restrictions of the) projections $(\Pi_*^{\text{gC}})_{(a, \phi)}$ and $\Pi_{(a, \phi)}^{\text{elC}}$ as inverse maps relating the enlarged local Coulomb slice to the global Coulomb slice:

$$(35) \quad \mathcal{K}_{(a, \phi)}^{\text{e}} \begin{array}{c} \xrightarrow{(\Pi_*^{\text{gC}})_{(a, \phi)}} \\ \xleftarrow{\Pi_{(a, \phi)}^{\text{elC}}} \end{array} \mathcal{T}_{(a, \phi)}^{\text{gC}}.$$

To see that they are inverse to each other, it suffices to note that both are given by adding a uniquely determined vector that is tangent to the \mathcal{G}° -orbit.

Here is the analogue of Lemma 3.1 for Π^{elC} instead of (Π_*^{gC}) .

Lemma 3.3. *Let $k \geq 2$. View the enlarged local Coulomb projection as a section of the bundle $\text{Hom}(TW, TC(Y))$ over W , i.e., a map from W to $\text{Hom}(W, \mathcal{C}(Y))$, given by*

$$(a, \phi) \mapsto ((b, \psi) \mapsto \Pi_{(a, \phi)}^{\text{elC}}(b, \psi)).$$

Then, this map extends to smooth maps between the Sobolev completions

$$W_k \rightarrow \text{Hom}(W_j, \mathcal{C}_j(Y))$$

for all $-k \leq j \leq k$.

Proof. The fact that the extension to Sobolev completions is well-defined was established in Lemma 3.2. Smoothness can be deduced from Lemma 3.1, using the fact that Π^{elC} (restricted to TW) is the inverse to Π_*^{gC} . \square

3.3. The Seiberg-Witten equations. Let $(Y, g, \mathfrak{s}, \mathbb{S})$ be as above. We let $\rho : TY \rightarrow \text{End}(\mathbb{S})$ be the Clifford multiplication. Further, for $a \in \Omega^1(Y; i\mathbb{R})$, we will use $D_a : \Gamma(\mathbb{S}) \rightarrow \Gamma(\mathbb{S})$ to denote the Dirac operator corresponding to the connection $A_0 + a$, and D for the case of $a = 0$.

Consider the Chern-Simons-Dirac (CSD) functional, \mathcal{L} , on $\mathcal{C}(Y)$:

$$\mathcal{L}(a, \phi) = \frac{1}{2} \left(\int_Y \langle \phi, D_a \phi \rangle - \int_Y a \wedge da \right).$$

We let \mathcal{X} denote the L^2 -gradient of CSD:

$$\mathcal{X}(a, \phi) = (*da + \tau(\phi, \phi), D_a\phi),$$

where $\tau(\phi, \phi) = \rho^{-1}(\phi\phi^*)_0$ is a quadratic function coming from the Clifford multiplication.

It is not difficult to check that the CSD functional is gauge-invariant (since $b_1(Y) = 0$). The critical points of \mathcal{L} are the solutions to the *Seiberg-Witten equations*,

$$\mathcal{X}(a, \phi) = 0.$$

If the spinor ϕ is identically 0, then the solution (a, ϕ) is said to be *reducible*.

By measuring the length of the enlarged local Coulomb projections of tangent vectors to W , we obtain a Riemannian metric \tilde{g} on W . Explicitly, for any tangent vector (b, ψ) to $(a, \phi) \in W$, we set⁵

$$(36) \quad \langle (b, \psi), (b', \psi') \rangle_{\tilde{g}} = \operatorname{Re} \langle \Pi_{(a, \phi)}^{\text{elC}}(b, \psi), \Pi_{(a, \phi)}^{\text{elC}}(b', \psi') \rangle_{L^2}.$$

For future reference, let us mention that, since $\Pi_{(a, \phi)}^{\text{elC}}$ is an L^2 -orthogonal projection, we can also write

$$(37) \quad \langle (b, \psi), (b', \psi') \rangle_{\tilde{g}} = \operatorname{Re} \langle (b, \psi), \Pi_{(a, \phi)}^{\text{elC}}(b', \psi') \rangle_{L^2}.$$

The metric \tilde{g} has the property that the trajectories of the gradient flow of \mathcal{L} restricted to W are precisely the global Coulomb projections of the original gradient flow trajectories in $\mathcal{C}(Y)$. Therefore, in the global Coulomb slice with the metric \tilde{g} , the (downward) gradient flow trajectories are given by

$$(38) \quad \frac{d}{dt}\gamma(t) = -(\Pi_*^{\text{gC}})_{\gamma(t)}\mathcal{X}(\gamma(t)),$$

where $\gamma(t) = (a(t), \phi(t))$. Note that $\mathcal{X} = \operatorname{grad} \mathcal{L}$ is perpendicular to the level sets of \mathcal{L} with respect to the L^2 metric, which contain the gauge orbits, so \mathcal{X} is automatically contained in the local Coulomb slices. We can split the right hand side of (38) into a linear part l and a nonlinear part c , and re-write the flow equation as

$$\frac{d}{dt}\gamma(t) = -(l + c)(\gamma(t)),$$

where

$$(39) \quad l(a, \phi) = (*da, D\phi)$$

$$(40) \quad c(a, \phi) = (\pi \circ \tau(\phi, \phi), \rho(a)\phi + \xi(\phi)\phi),$$

with $\xi(\phi) : Y \rightarrow i\mathbb{R}$ being characterized by $d\xi(\phi) = (1 - \pi) \circ \tau(\phi, \phi)$ and $\int_Y \xi(\phi) = 0$. Recall also that π denotes the orthogonal projection to $\ker d^*$.

The \tilde{g} -gradient of the restriction $\mathcal{L}|_W$ extends to a map

$$\mathcal{X}^{\text{gC}} = l + c : W_k \rightarrow W_{k-1},$$

such that l is a linear Fredholm operator, and c is quadratic.⁶ The linear operator l is self-adjoint with respect to the L^2 inner product, but not necessarily \tilde{g} . Further, the map c is continuous as a map from W_k to W_k , and thus compact as a map from W_k to W_{k-1} . The corresponding flow lines are called *Seiberg-Witten trajectories* (in Coulomb gauge). Such a

⁵The L^2 inner product is Hermitian, and thus the real part gives a real inner product. This is what we need to define a Riemannian metric.

⁶We chose our conventions to be in agreement with [KM07]. In [Man03], the map $l + c$ went from W_{k+1} to W_k .

trajectory $\gamma = (a(t), \phi(t)) : \mathbb{R} \rightarrow W$ is said to be *of finite type* if $\mathcal{L}(\gamma(t))$ and $\|\phi(t)\|_{C^0}$ are bounded in t .

3.4. Finite-dimensional approximation. For $\lambda > 1$, let us denote by W^λ the finite-dimensional subspace of W spanned by the eigenvectors of l with eigenvalues in the interval $(-\lambda, \lambda)$.⁷ The L^2 orthogonal projection from W to W^λ will be denoted \tilde{p}^λ . We modify this to make it smooth in λ , using the following preliminary definition:

$$(41) \quad p_{\text{prel}}^\lambda = \int_0^1 \beta(\theta) \tilde{p}_{-\lambda+\theta}^{\lambda-\theta} d\theta,$$

where β is a smooth, non-negative function that is non-zero exactly on $(0, 1)$, and such that $\int_{\mathbb{R}} \beta(\theta) d\theta = 1$.

Our p_{prel}^λ was the one used in [Man03], where it was denoted p^λ . In this paper, it will be convenient to arrange for the smoothed projection to be the actual projection \tilde{p}^λ at an infinite sequence of λ 's. Let us fix such a sequence:

$$\lambda_1^\bullet < \lambda_2^\bullet < \dots$$

such that $\lambda_i^\bullet \rightarrow \infty$ and none of the λ_i are eigenvalues of l . Fix also disjoint intervals $[\lambda_i^\bullet - \epsilon_i, \lambda_i^\bullet + \epsilon_i]$ that do not contain eigenvalues of l . Choose smooth bump functions $\beta_i : (0, \infty) \rightarrow [0, 1]$ supported in $[\lambda_i^\bullet - \epsilon_i, \lambda_i^\bullet + \epsilon_i]$ and with $\beta_i(\lambda_i^\bullet) = 1$. Then set

$$(42) \quad p^\lambda = \sum_i \beta_i(\lambda) \tilde{p}^\lambda + (1 - \sum_i \beta_i(\lambda)) p_{\text{prel}}^\lambda.$$

We now have that p^λ is smooth in λ , and $p^\lambda = \tilde{p}^\lambda$ for $\lambda \in \{\lambda_1^\bullet, \lambda_2^\bullet, \dots\}$. Moreover, for any λ , the image of p^λ is the subspace W^λ .

As an aside, let us remark that in [Man03] there is a different definition of the L_k^2 Sobolev norm on W , using l instead of the covariant derivative ∇ . This produces a norm equivalent to the usual one. With the definition in [Man03], the L^2 projection \tilde{p}^λ would have been the orthogonal projection to W^λ with respect to the L_k^2 metric for any k , and we would have had $\|p^\lambda(x)\|_{L_k^2} \leq \|x\|_{L_k^2}$. Given our choice of Sobolev norms, the same inequality holds up to a constant:

$$\|p^\lambda(x)\|_{L_k^2} \leq \Theta_k \|x\|_{L_k^2},$$

where Θ_k depends only on k and the Riemannian manifold Y . We choose $\Theta_0 = 1$.

On W^λ , we consider the flow equation

$$(43) \quad \frac{d}{dt} \gamma(t) = -(l + p^\lambda c)(\gamma(t)).$$

We refer to solutions of (43) as *approximate Seiberg-Witten trajectories*.

⁷In [Man03], the role of W^λ was played by W_λ^μ , the subspace spanned by eigenvectors with eigenvalues between λ and μ . In this paper we restrict to $\mu = -\lambda$; this produces the same spectrum. The reason for the change is that in Section 7.2 we will need to turn the parameter space for λ into a manifold with boundary, and it is easier to do so with only one degree of freedom. Also, notice that in [Man03, Section 4], the approximation W_λ^μ was initially defined as the span of eigenvectors with eigenvalues in $(\lambda, \mu]$, but later changed so that it is the image of the smoothed projection p_λ^μ . This means using the open interval (λ, μ) instead of $(\lambda, \mu]$, and in fact it is easier to define the approximation as such from the beginning. In our setting we use $(-\lambda, \lambda)$.

Remark 3.4. If we consider the restriction of the CSD functional to W^λ , its gradient with respect to \tilde{g} is

$$\tilde{p}_g^\lambda(l+c)\tilde{p}_g^\lambda = l + \tilde{p}_g^\lambda c \tilde{p}_g^\lambda,$$

where \tilde{p}_g^λ denotes the \tilde{g} -orthogonal projection onto W^λ . It would be rather cumbersome to work with these projections, so we replaced \tilde{p}_g^λ with the L^2 orthogonal projection \tilde{p}^λ . When λ is one of the cut-offs λ_i^\bullet , on W^λ we have

$$\tilde{p}^\lambda(l+c)\tilde{p}^\lambda = l + \tilde{p}^\lambda c = l + p^\lambda c.$$

However, even in this case (when $\lambda = \lambda_i^\bullet$), we expect that $l + p^\lambda c$ is neither the L^2 nor the \tilde{g} gradient of a function. We will show in Section 8 that $l + p^\lambda c$ is a quasi-gradient, so it can still be used to do Morse theory.

Fix a natural number $k \geq 5$. There exists a constant $R > 0$, such that all Seiberg-Witten trajectories $\gamma : \mathbb{R} \rightarrow W$ of finite type are contained in $B(R)$, the ball of radius R in W_k . The following is a corresponding compactness result for approximate Seiberg-Witten trajectories:

Proposition 3.5 (Proposition 3 in [Man03]). *For any λ sufficiently large (compared to R), if $\gamma : \mathbb{R} \rightarrow W^\lambda$ is a trajectory of $(l + p^\lambda c)$, and $\gamma(t)$ is in $\overline{B(2R)}$ for all t , then in fact $\gamma(t)$ is contained in $B(R)$.*

This was proved in [Man03] with p_{prel}^λ instead of p^λ , but the same arguments work in our setting.

3.5. The Conley index and the Seiberg-Witten Floer spectrum. Recall the definition of the Conley index from Section 2.4, and of the S^1 -equivariant refinement I_{S^1} mentioned at the end of Section 2.8.

With this in mind, we are ready to define the Seiberg-Witten Floer spectrum. We fix k , R , and sufficiently large λ such that Proposition 3.5 applies. We consider the vector field $u^\lambda(l + p^\lambda c)$ on W^λ , where u^λ is a smooth, S^1 -invariant, cut-off function on W^λ that vanishes outside of $B(3R)$. This generates the flow ϕ^λ that we will work with. Denote by S^λ the union of all trajectories of ϕ^λ inside $\overline{B(R)}$. Recall from Proposition 3.5 that these are the same as the trajectories that stay in $\overline{B(2R)}$. This implies that S^λ is an isolated invariant set.

Since everything is S^1 -invariant, we can construct the equivariant Conley index $I^\lambda = I_{S^1}(\phi^\lambda, S^\lambda)$. We must de-suspend appropriately to make the stable homotopy type independent of λ :

$$\text{SWF}(Y, \mathfrak{s}, g) = \Sigma^{-W^{(-\lambda, 0)}} I^\lambda,$$

where $W^{(-\lambda, 0)}$ denotes the direct sum of the eigenspaces of l with eigenvalues in the interval $(-\lambda, 0)$. As we vary the metric g , the spectrum $\text{SWF}(Y, \mathfrak{s}, g)$ varies by suspending (or de-suspending) with copies of the vector space \mathbb{C} . In [Man03], this indeterminacy is fixed by introducing a quantity $n(Y, \mathfrak{s}, g) \in \mathbb{Q}$ (a linear combination of eta invariants), and setting

$$\text{SWF}(Y, \mathfrak{s}) = \Sigma^{-n(Y, \mathfrak{s}, g)} \text{SWF}(Y, \mathfrak{s}, g),$$

where the de-suspension by rational numbers is defined formally. For the definition of $n(Y, \mathfrak{s}, g)$, see (233) below. We have:

Theorem 3.6 (Theorem 1 in [Man03]). *The S^1 -equivariant stable homotopy type of $\text{SWF}(Y, \mathfrak{s})$ is an invariant of the pair (Y, \mathfrak{s}) .*

Remark 3.7. The construction of $\text{SWF}(Y, \mathfrak{s})$ in [Man03] used the smoothed projections p_{prel}^λ . By interpolating linearly between p_{prel}^λ and p^λ , and using the homotopy invariance properties of the Conley index, we see that the definitions from p^λ and p_{prel}^λ yield the same $\text{SWF}(Y, \mathfrak{s})$.

4. MONOPOLE FLOER HOMOLOGY

In this section we review the definition of monopole Floer homology given by Kronheimer and Mrowka in their book [KM07]. Recall that we are considering the case of a rational homology sphere Y and Spin^c structure \mathfrak{s} , which is necessarily torsion. In this case, all of the reducible solutions to the Seiberg-Witten equations are gauge equivalent. While we worked in the Coulomb gauge for the Floer spectrum, for now, we will return to the entire configuration space

$$\mathcal{C}(Y) = \Omega^1(Y; i\mathbb{R}) \oplus \Gamma(\mathbb{S})$$

(where we have made this identification via a fixed flat Spin^c connection).

Here are the main ideas in the construction of monopole Floer homology: One would like to proceed by analogy with Morse homology; that is, to build a chain complex whose generators are given by gauge-equivalence classes of critical points of the CSD functional and whose differential counts gradient trajectories between them. However, there are several technical issues that need to be addressed. One issue is that the gauge group does not act freely near the reducible solutions to the Seiberg-Witten equations, and these points cause a serious problem. This problem was solved by Kronheimer and Mrowka by blowing up the singular set, in a way similar to the construction of S^1 -equivariant Morse homology in Section 2.6. The second problem is that even after performing a blow-up, transversality may still not be satisfied for the moduli spaces. For this reason, special perturbations to the CSD functional must be introduced.

4.1. Seiberg-Witten equations on the blow-up. We will often work on cylinders of the form $Z = I \times Y$, where I is an interval (possibly \mathbb{R}). The Spin^c structure on Y induces a Spin^c structure on Z with unitary rank 4 bundle $\mathbb{S}^+ \oplus \mathbb{S}^-$ on Z ; here, \mathbb{S}^\pm are the ∓ 1 eigenspaces of $\rho(d \text{vol}_Z)$, where ρ is the induced Clifford multiplication for differential forms on Z . Both \mathbb{S}^+ and \mathbb{S}^- can be identified with the pull-backs of \mathbb{S} under the projection $p : Z \rightarrow Y$. Again identifying 1-forms with Spin^c connections, we have

$$\mathcal{C}(Z) = \{(a, \phi) \mid a \in \Omega^1(Z; i\mathbb{R}), \phi \in \Gamma(\mathbb{S}^+)\}.$$

This is acted on by the gauge group $\mathcal{G}(Z) = C^\infty(Z, S^1)$. An element of $\mathcal{C}(Z)$ in temporal gauge (i.e., such that the dt component of a is zero) corresponds to a path $\gamma(t) = (a(t), \phi(t))$ in $\mathcal{C}(Y)$, and we will often not distinguish between these. Note that every element of $\mathcal{C}(Z)$ is gauge-equivalent to a configuration in temporal gauge.

We write down the four-dimensional Seiberg-Witten equations on Z as $\mathcal{F}(a, \phi) = 0$, where

$$\mathcal{F}(a, \phi) = (d^+a - \rho^{-1}((\phi\phi^*)_0), D_a^+\phi) \in \Omega_+^2(Z; i\mathbb{R}) \oplus \Gamma(\mathbb{S}^-),$$

and $D_a^+ : \Gamma(\mathbb{S}^+) \rightarrow \Gamma(\mathbb{S}^-)$ is the Dirac operator and d^+a is the self-dual part of da . If $\mathcal{F}(\gamma) = 0$, then γ represents a *Seiberg-Witten trajectory*. If γ is the constant path, then this is giving a solution to the Seiberg-Witten equations on Y . In general, $\mathcal{F}(\gamma) = 0$ is equivalent to γ being a downward gradient trajectory of \mathcal{L} . We will often go back and forth between this notion of trajectories of $\text{grad } \mathcal{L} = \mathcal{X}$ on Y and solutions to the Seiberg-Witten equations on Z . We will also think of \mathcal{F} as a section of $\mathcal{C}(Z)$ into $\mathcal{V}(Z)$, the trivial $\Gamma(Z; i\Lambda_+^2 T^*Z \oplus \mathbb{S}^-)$ bundle over $\mathcal{C}(Z)$.

As discussed above, in order to deal with the reducible solution, we must blow-up our configuration spaces. We first consider

$$\mathcal{C}^\sigma(Y) = \{(a, s, \phi) \mid s \geq 0, \|\phi\|_{L^2} = 1\} \subset \Omega^1(Y; i\mathbb{R}) \times \mathbb{R}_{\geq 0} \times \Gamma(\mathbb{S}).$$

We can define $\mathcal{C}^\sigma(Z)$ similarly, as the space of triples $(a, s, \phi) \in \Omega^1(Z; i\mathbb{R}) \times \mathbb{R}_{\geq 0} \times \Gamma(\mathbb{S}^+)$ with $\|\phi\|_{L^2(Z)} = 1$. The *blown-up Seiberg-Witten equations* on $\mathcal{C}^\sigma(Z)$ are given by $\mathcal{F}^\sigma(a, s, \phi) = 0$, where \mathcal{F}^σ is a section of a bundle $\mathcal{V}^\sigma(Z)$ over $\mathcal{C}^\sigma(Z)$, defined to be the pullback of $\mathcal{V}(Z)$ under the blow-down from $\mathcal{C}^\sigma(Z)$ to $\mathcal{C}(Z)$; see [KM07, p.115]. Explicitly, we have

$$\mathcal{F}^\sigma(a, s, \phi) = (d^+a - s^2\rho^{-1}((\phi\phi^*)_0), D_a^+\phi).$$

However, there is a variant of the four-dimensional blow-up that is more directly related to paths in $\mathcal{C}^\sigma(Y)$. This is the so-called τ *model* defined in [KM07, Section 6.3]. Precisely, we let

$$\mathcal{C}^\tau(Z) \subset \Omega^1(Z; i\mathbb{R}) \times C^\infty(I) \times C^\infty(Z; \mathbb{S}^+)$$

be the space of triples (a, s, ϕ) with $s(t) \geq 0$ and $\|\phi(t)\|_{L^2(Y)} = 1$ for all $t \in I$. After moving it into temporal gauge, an element of $\mathcal{C}^\tau(Z)$ determines a path $\gamma^\sigma(t) = (a(t), s(t), \phi(t))$ in $\mathcal{C}^\sigma(Y)$.

In temporal gauge on $\mathcal{C}^\tau(Z)$, the blown-up Seiberg-Witten equations can be written as

$$\begin{aligned} \frac{d}{dt}a(t) &= - * da(t) - s(t)^2\tau(\phi(t), \phi(t)) \\ \frac{d}{dt}s(t) &= -\Lambda(a(t), s(t), \phi(t))s(t) \\ \frac{d}{dt}\phi(t) &= -D_{a(t)}\phi(t) + \Lambda(a(t), s(t), \phi(t))\phi(t), \end{aligned}$$

where

$$(44) \quad \Lambda(a, s, \phi) = \langle \phi, D_a\phi \rangle_{L^2}.$$

Furthermore, the right hand sides of these equations induce a vector field, denoted \mathcal{X}^σ , on $\mathcal{C}^\sigma(Y)$ whose flow trajectories are precisely the solutions to the three equations above. The zeros of \mathcal{X}^σ on $\mathcal{C}^\sigma(Y)$ are easily rephrased in terms of the zeros of \mathcal{X} on $\mathcal{C}(Y)$:

Proposition 4.1 (Proposition 6.2.3 in [KM07]). *If $s > 0$, then $(a, s, \phi) \in \mathcal{C}^\sigma(Y)$ is a zero of \mathcal{X}^σ in $\mathcal{C}^\sigma(Y)$ if and only if $(a, s\phi)$ is a zero of \mathcal{X} in $\mathcal{C}(Y)$. If $s = 0$, then (a, s, ϕ) is a zero if and only if $(a, 0)$ is a zero of \mathcal{X} and ϕ is an eigenvector of D_a .*

Remark 4.2. Since $b_1(Y) = 0$, we have that $(a, 0)$ is a zero of \mathcal{X} if and only if $a = 0$.

Throughout this section, we fix an integer $k \geq 2$. We may also blow up the Sobolev completions to obtain spaces $\mathcal{C}_k^\sigma(Y)$ such that \mathcal{X}^σ extends. Similarly, we can complete the gauge group \mathcal{G} to \mathcal{G}_k . This leads to the various blown-up configuration spaces mod gauge:

$$\mathcal{B}(Y) = \mathcal{C}(Y)/\mathcal{G}, \quad \mathcal{B}^\sigma(Y) = \mathcal{C}^\sigma(Y)/\mathcal{G}, \quad \mathcal{B}_k(Y) = \mathcal{C}_k(Y)/\mathcal{G}_{k+1}, \quad \mathcal{B}_k^\sigma(Y) = \mathcal{C}_k^\sigma(Y)/\mathcal{G}_{k+1}.$$

Recall that for $j \leq k$, \mathcal{T}_j denotes the L_j^2 completion of $T\mathcal{C}_k(Y)$. We have the analogous construction in the blow-up, \mathcal{T}_j^σ , which decomposes as

$$(45) \quad \mathcal{T}_j^\sigma = \mathcal{J}_j^\sigma \oplus \mathcal{K}_j^\sigma,$$

where \mathcal{J}_k^σ consists of the tangents to the gauge orbits. More explicitly, the L_j^2 completion of the tangent space at $x = (a, s, \phi) \in \mathcal{C}_k^\sigma(Y)$ is

$$\mathcal{T}_{j,x}^\sigma = \{(b, r, \psi) \mid \operatorname{Re}\langle \phi, \psi \rangle_{L^2} = 0\} \subset L_j^2(Y; iT^*Y) \oplus \mathbb{R} \oplus L_j^2(Y; \mathbb{S}).$$

At $\phi = 0$, we also want to consider $\mathcal{T}_j^{\operatorname{red}}$, the completion of the tangent bundle to $L_k^2(\Omega^1(Y; i\mathbb{R}))$. We have an analogous splitting of $\mathcal{T}_{k,a}^{\operatorname{red}}$ into $\mathcal{J}_{k,a}^{\operatorname{red}} = L_k^2(\operatorname{im} d)$ and $\mathcal{K}_{k,a}^{\operatorname{red}} = L_k^2(\ker d^*)$.

In four dimensions, we can similarly divide the blown up configuration space $\mathcal{C}^\sigma(Z)$ by gauge to obtain a space $\mathcal{B}^\sigma(Z)$. In the τ model, we define the quotient configuration space $\mathcal{B}^\tau(Z) = \mathcal{C}^\tau(Z)/\mathcal{G}(Z)$. Also, starting from the bundle $\mathcal{V}(Z)$ over $\mathcal{C}(Z)$, we obtain a bundle $\mathcal{V}^\tau(Z)$ over $\mathcal{C}^\tau(Z)$. Explicitly, the fiber of $\mathcal{V}^\tau(Z)$ over (a, s, ϕ) is

$$\begin{aligned} \mathcal{V}^\tau(Z)_{(a,s,\phi)} &= \{(b, r, \psi) \mid \operatorname{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0, \forall t\} \\ &\subset C^\infty(Z; i\Lambda_+^2 T^*Z) \oplus C^\infty(\mathbb{R}) \oplus C^\infty(Z; \mathbb{S}^-). \end{aligned}$$

In the τ model, the blown-up Seiberg-Witten equations on Z can be written as the zeros of a section \mathcal{F}^τ of the bundle $\mathcal{V}^\tau(Z)$; see [KM07, Equations (6.11)] for the exact formula. In temporal gauge, we simply have

$$\mathcal{F}^\tau = \frac{d}{dt} + \mathcal{X}^\sigma.$$

When Z is a compact cylinder, that is, $Z = [t_1, t_2] \times Y$, we define the Sobolev completions $\mathcal{C}_k(Z)$, $\mathcal{B}_k(Z)$, $\mathcal{V}_k(Z)$, and similarly with σ or τ superscripts. We also have completed tangent bundles $\mathcal{T}_j(Z)$, $\mathcal{T}_j^\sigma(Z)$ and $\mathcal{T}_j^\tau(Z)$. We should note that $\mathcal{C}_k^\sigma(Y)$ and $\mathcal{C}_k^\sigma(Z)$ are Hilbert manifolds with boundary, so the tangent bundles make sense. The τ model $\mathcal{C}_k^\tau(Z)$ is not a Hilbert manifold, even with boundary. Nevertheless, it is a closed subset of the Hilbert manifold $\tilde{\mathcal{C}}_k^\tau(Z)$, which is the L_k^2 completion of

$$(46) \quad \tilde{\mathcal{C}}^\tau(Z) = \{(a, s, \phi) \mid \|\phi(t)\|_{L^2(Y)} = 1, \forall t\} \subset \Omega^1(Z; i\mathbb{R}) \times C^\infty([t_1, t_2]) \times C^\infty(Z; \mathbb{S}^+);$$

see [KM07, Section 9.2]. By the completed tangent bundle to $\mathcal{C}_k^\tau(Z)$ we simply mean the restriction of the completed tangent bundle to $\tilde{\mathcal{C}}_k^\tau(Z)$, i.e. the L_j^2 completion of the bundle $\mathcal{T}^\tau(Z)$ with fibers

$$\begin{aligned} \mathcal{T}^\tau(Z)_{(a,s,\phi)} &= \{(b, r, \psi) \mid \operatorname{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0, \forall t\} \\ &\subset C^\infty(Z; iT^*Z) \oplus C^\infty([t_1, t_2]) \oplus C^\infty(Z; \mathbb{S}^-). \end{aligned}$$

The bundle $\mathcal{V}^\tau(Z)$ extends naturally over $\tilde{\mathcal{C}}^\tau(Z)$ as well. When discussing $\mathcal{V}^\tau(Z)$, this will refer to the extension and will explicitly mention the restriction to $\mathcal{C}^\tau(Z)$ explicitly when it is used. Finally, we will need to quotient the extended spaces $\tilde{\mathcal{C}}_k^\tau(Z)$ by gauge, which we write as $\tilde{\mathcal{B}}_k^\tau$. When Z is non-compact, for example $Z = \mathbb{R} \times Y$, it is often more useful to consider the local Sobolev completions $L_{k,loc}^2$; see [KM07, Section 13.1] for more details.

As mentioned in Section 3, gauge transformations preserve the property of being a Seiberg-Witten trajectory; this extends through blow-ups and Sobolev completions. Monopole Floer homology will be defined as the Morse homology (for manifolds with boundary) of a perturbation of the Seiberg-Witten equations on $\mathcal{B}_k^\sigma(Y)$.

4.2. Perturbed Seiberg-Witten equations. Consider a function $f : \mathcal{C}(Y) \rightarrow \mathbb{R}$ which is gauge-invariant. A *perturbation* is a section $\mathfrak{q} : \mathcal{C}(Y) \rightarrow \mathcal{T}_0$. We will call \mathfrak{q} the *formal gradient* of f , or $\operatorname{grad} f$ (even though it's not actually a gradient), if for all $\gamma \in C^\infty([0, 1], \mathcal{C}(Y))$

$$\int_0^1 \left\langle \frac{d}{dt} \gamma(t), \mathfrak{q}(\gamma(t)) \right\rangle_{L^2} dt = f \circ \gamma(1) - f \circ \gamma(0).$$

We will use the notation

$$\mathcal{X}_{\mathfrak{q}} := \mathcal{X} + \mathfrak{q} = (\operatorname{grad} \mathcal{L}) + \mathfrak{q}.$$

We will also write \mathcal{L}_q for $\mathcal{L} + f$, so $\mathcal{X}_q = \text{grad } \mathcal{L}_q$. Note that q is described by its components q^0 to $\Omega^1(Y; i\mathbb{R})$ and q^1 to $\Gamma(\mathbb{S})$. More generally, from now on we will use the superscripts 0 and 1 to describe the 1-form and spinorial parts of vectors in $\mathcal{C}(Y)$.

The perturbation q also induces a section $\hat{q} : \mathcal{C}(Z) \rightarrow \mathcal{V}_0(Z)$. When $Z = [t_1, t_2] \times Y$ is a compact cylinder, we will require that this actually extends to a smooth section $\mathcal{C}_k(Z) \rightarrow \mathcal{V}_k(Z)$.

The flow trajectories $(a(t), \phi(t))$ of \mathcal{X}_q are solutions to the equations

$$\begin{aligned} \frac{d}{dt}a &= - * da - \tau(\phi, \phi) - q^0(a, \phi) \\ \frac{d}{dt}\phi &= -D_a\phi - q^1(a, \phi). \end{aligned}$$

Recast on Z , these are the *perturbed Seiberg-Witten equations*, $\mathcal{F}_q = \mathcal{F} + \hat{q} = 0$.

We now move to the blow-up. The vector field \mathcal{X}_q induces a vector field \mathcal{X}_q^σ on $\mathcal{C}^\sigma(Y)$. Precisely, on $\mathcal{C}^\sigma(Y)$, we define

$$q^\sigma(a, s, \phi) = (q^0(a, s\phi), \text{Re}\langle \tilde{q}^1(a, s, \phi), \phi \rangle_{L^2} \cdot s, \tilde{q}^1(a, s, \phi)^\perp),$$

where $\tilde{q}^1(a, s, \phi) = \int_0^1 \mathcal{D}_{(a, st\phi)} q^1(0, \phi) dt$ and \perp means projection onto the real orthogonal complement of ϕ . We can obtain \hat{q}^σ similarly, with components $(q^{\hat{\sigma}, 0}, q^{\hat{\sigma}, 1})$. This leads to the flow equations for \mathcal{X}_q^σ :

$$\begin{aligned} \frac{d}{dt}a &= - * da - s^2 \tau(\phi, \phi) - q^0(a, s\phi), \\ \frac{d}{dt}s &= -\Lambda_q(a, s, \phi)s, \\ \frac{d}{dt}\phi &= -D_a\phi - \tilde{q}^1(a, s, \phi) + \Lambda_q(a, s, \phi)\phi, \end{aligned}$$

where

$$(47) \quad \Lambda_q(a, s, \phi) = \text{Re}\langle \phi, D_a\phi + \tilde{q}^1(a, s, \phi) \rangle_{L^2}.$$

These are the solutions to the *perturbed Seiberg-Witten equations on the blow-up* in temporal gauge, or $\mathcal{F}_q^\sigma = 0$ for short. We call $\Lambda_q(a, s, \phi)$ the *spinorial energy* of (a, s, ϕ) . Note that for an irreducible stationary point (a, s, ϕ) , we must have $\Lambda_q(a, s, \phi) = 0$.

We can now state the perturbed analogue of Proposition 4.1. First, we define the operator

$$D_{q,a} : L_k^2(Y; \mathbb{S}) \rightarrow L_{k-1}^2(Y; \mathbb{S}), \quad D_{q,a}(\phi) = D_a\phi + \mathcal{D}_{(a,0)} q^1(0, \phi).$$

Proposition 4.3 (Proposition 10.3.1 of [KM07]). *The element (a, s, ϕ) in $\mathcal{C}_k^\sigma(Y)$ is a zero of \mathcal{X}_q^σ if and only if:*

- (a) $s \neq 0$ and $(a, s\phi)$ is a stationary point of \mathcal{X}_q , or
- (b) $s = 0$ and $(a, 0)$ is a stationary point of \mathcal{X}_q and ϕ is an eigenvector of $D_{q,a}$.

Remark 4.4. A stationary point of the form $(a, 0)$ is called *reducible*. Recall from Remark 4.2 that the unperturbed Seiberg-Witten vector field \mathcal{X} has a unique reducible stationary point mod gauge, namely $a = 0$. We claim that, when q is small, the perturbed field \mathcal{X}_q also has a unique reducible stationary point. Indeed, notice that the linear map $*d : L_k^2(Y; iT^*Y) \rightarrow L_{k-1}^2(Y; iT^*Y)$ induces an invertible map from $(\ker d^*)_k$ to $(\ker d^*)_{k-1}$. Note also that $(\ker d^*)_j$ is the L^2 orthogonal complement in $L_j^2(Y; iT^*Y)$ of the tangents to the gauge orbits of connections. Therefore, by the inverse function theorem, for a small perturbation q , there

is a unique $a \in L_k^2(Y; iT^*Y)$ mod gauge, in a neighborhood of 0, satisfying $*da + \mathfrak{q}^0(a, 0) = 0$. (Here we are using that \mathfrak{q} is L^2 orthogonal to the tangents of the gauge-orbits, since it is the formal gradient of a gauge-equivariant function.) By elliptic bootstrapping, a is smooth. Further, by gauge-equivariance, we have $\mathfrak{q}^1(a, 0) = 0$ and we can conclude that $(a, 0)$ provides the unique reducible solution to $\mathcal{X}_{\mathfrak{q}}$ near zero. Using the compactness properties of the perturbed Seiberg-Witten equations, it also follows that (for sufficiently small \mathfrak{q}) there are no reducible solutions outside the fixed neighborhood of 0.

We will not, however, use the uniqueness of reducible solutions for small perturbations in this paper.

The (perturbed, blown-up) Seiberg-Witten map $\mathcal{F}_{\mathfrak{q}}^{\sigma}$ can be viewed as a section of the bundle $\mathcal{V}^{\sigma}(Z)$ over $\mathcal{C}^{\sigma}(Z)$. In the τ model, there is a similar section $\mathcal{F}_{\mathfrak{q}}^{\tau}$ of $\mathcal{V}^{\tau}(Z)$ over $\tilde{\mathcal{C}}^{\tau}(Z)$. In temporal gauge, we can simply write

$$\mathcal{F}_{\mathfrak{q}}^{\tau} = \frac{d}{dt} + \mathcal{X}_{\mathfrak{q}}^{\sigma}.$$

4.3. Tame perturbations. We are interested in studying the moduli spaces of flows of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ connecting stationary points. In order to obtain the desired compactification results for the moduli spaces, we require some conditions on \mathfrak{q} .

Let $Z = [t_1, t_2] \times Y$ be a compact cylinder. For every one-form $a \in \Omega^1(Z; i\mathbb{R})$, we can consider the Sobolev norm $L_{k,a}^2$ defined using as covariant derivative the connection corresponding to a , namely ∇_{A_0+a} . The $L_{k,a}^2$ norm is equivalent to the usual Sobolev norm, and gives rise to the same Sobolev completion. However, when we want to state global bounds it becomes important to specify the precise Sobolev norm. We will take the usual norm on $\mathcal{C}_k(Y), \mathcal{C}_k(Z)$. However, on the bundles $\mathcal{V}_k(Z)$ and $T\mathcal{C}_k(Z)$, in the fiber over (a, ϕ) we will take the $L_{k,a}^2$ norm; this turns them into gauge-invariant normed vector bundles.

Definition 4.5 (Definition 10.5.1 in [KM07]). *Fix an integer $k \geq 2$. Suppose that a section $\mathfrak{q} : \mathcal{C}(Y) \rightarrow \mathcal{T}_0$ is the formal gradient of a continuous, \mathcal{G} -invariant function. We say \mathfrak{q} is a k -tame perturbation if, for any compact cylinder $Z = [t_1, t_2] \times Y$, the following hold:*

- (i) $\hat{\mathfrak{q}}$ defines an element of $C^{\infty}(\mathcal{C}_k(Z), \mathcal{V}_k(Z))$;
- (ii) $\hat{\mathfrak{q}}$ also defines an element of $C^0(\mathcal{C}_j(Z), \mathcal{V}_j(Z))$, for all integers $j \in [1, k]$;
- (iii) The derivative $\mathcal{D}\hat{\mathfrak{q}} \in C^{\infty}(\mathcal{C}_k(Z), \text{Hom}(\mathcal{T}_k(Z), \mathcal{V}_k(Z)))$ extends to a smooth map into $\text{Hom}(\mathcal{T}_j(Z), \mathcal{V}_j(Z))$, for all integers $j \in [-k, k]$;
- (iv) There exists a constant m such that

$$\|\mathfrak{q}(a, \phi)\|_{L^2} \leq m(\|\phi\|_{L^2} + 1),$$

for all $(a, \phi) \in \mathcal{C}_k(Y)$;

- (v) For any $a_0 \in i\Omega^1(Z)$, there is a function $h : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$\|\hat{\mathfrak{q}}(a, \phi)\|_{L_{1,a}^2} \leq h(\|(a, \phi)\|_{L_{1,a_0}^2}),$$

for all $(a, \phi) \in \mathcal{C}_k(Z)$;

- (vi) \mathfrak{q} extends to a C^1 section of $\mathcal{C}_1(Y)$ into \mathcal{T}_0 .

We say that \mathfrak{q} is tame if it is k -tame for all $k \geq 2$.

Remark 4.6. Conditions (i), (ii), (iii) and (v) are phrased in terms of the four-dimensional perturbation $\hat{\mathfrak{q}}$, but they readily imply the analogous statements for \mathfrak{q} . For example, from part (i) in Definition 4.5, we know that for a tame perturbation \mathfrak{q} , $\hat{\mathfrak{q}}$ gives a smooth map from $\mathcal{C}_k(Z)$ to $\mathcal{V}_k(Z)$, for all $k \geq 1$. For $x \in \mathcal{C}_k(Y)$, we can consider the element of $\mathcal{C}_k(Z)$

that is constantly x in every slice $\{t\} \times Y$. Its image under $\hat{\mathbf{q}}$ is constantly $\mathbf{q}(x)$ in every slice, so we conclude that $\mathbf{q}(x)$ is in $\mathcal{V}_k(Y)$.

For a tame perturbation \mathbf{q} , the set of solutions to $\mathcal{X}_{\mathbf{q}} = 0$, topologized as a subspace of $\mathcal{B}_k(Y)$, is compact for tame \mathbf{q} . In this case, any trajectory of $\mathcal{X}_{\mathbf{q}}^\sigma$ between two stationary points in $\mathcal{C}_k^\sigma(Y)$ is gauge-equivalent to a smooth trajectory, which thus lives in $\mathcal{C}^\sigma(Y)$; no Sobolev completion is necessary.

Kronheimer and Mrowka construct explicitly a collection of tame perturbations which are the gradient of cylinder functions. We will not define these here. However, we will always need to work with a space of perturbations that does contain them.

Definition 4.7. *A separable Banach space \mathcal{P} is a large Banach space of tame perturbations if there exists a map $\mathcal{P} \rightarrow C^0(\mathcal{C}(Y), \mathcal{T}_0)$ which takes x to \mathbf{q}_x such that:*

- (i) *the image of this map contains a countably infinite family of perturbations which are gradients of cylinder functions,*
- (ii) *\mathbf{q}_x is tame for all $x \in \mathcal{P}$,*
- (iii) *for $k \geq 2$, the map $\mathcal{P} \times \mathcal{C}_k(I \times Y) \rightarrow \mathcal{V}_k(I \times Y)$ is smooth for any compact interval I in \mathbb{R} ,*
- (iv) *the map $\mathcal{P} \times \mathcal{C}_1(Y) \rightarrow \mathcal{T}_1(Y)$ is continuous, and*
- (v) *there exist a constant m and a map $\mu : \mathbb{R} \rightarrow \mathbb{R}$ such that $\mathcal{P} \times \mathcal{C}_1(Y) \rightarrow \mathcal{T}_1(Y)$ satisfies*

$$\begin{aligned} \|\mathbf{q}_x(a, \phi)\|_{L^2} &\leq m\|x\|(\|\phi\|_{L^2} + 1) \\ \|\mathbf{q}_x(a, \phi)\|_{L_{1,a_0}^2} &\leq \|x\|\mu(\|(a, \phi)\|_{L_{1,a_0}^2}). \end{aligned}$$

4.4. Very tame perturbations. We now make a digression. For the purposes of doing finite dimensional approximation (as we shall do in this paper), we need slightly stronger assumptions on the perturbation \mathbf{q} than the ones in Definition 4.5. We give the relevant definitions here. These concepts are new material; they did not appear in [KM07].

Let us recall that a linear operator between Banach spaces, $f : E \rightarrow F$, is continuous if and only if it is bounded, i.e., there is a constant K such that $\|f(x)\| \leq K\|x\|$ for all x . This condition is also equivalent to requiring f to take bounded sets to bounded sets. For nonlinear operators, we introduce the following terminology.

Definition 4.8. *Let E and F be Banach spaces. A map $f : E \rightarrow F$ is called functionally bounded if $f(B) \subset F$ is bounded whenever $B \subset E$ is bounded. In other words, there exists a function $h : \mathbb{R} \rightarrow \mathbb{R}$ such that*

$$\|f(x)\| \leq h(\|x\|),$$

for all $x \in E$.

We use the term *functionally bounded*, rather than *bounded*, to prevent confusion with the usual notion of bounded functions in analysis, which requires that $\|f(x)\| \leq K$ for some K .

For non-linear operators, continuity does not imply functional boundedness; see [Exc]. We will need both of these conditions, so let us denote by

$$C_{\text{fb}}^0(E, F)$$

the space of continuous and functionally bounded operators from E to F . Further, for $m \geq 1$ (and also for $m = \infty$), we let

$$C_{\text{fb}}^m(E, F) = \{f \in C^m(E, F) \mid \mathcal{D}^\ell f \in C_{\text{fb}}^0(E, \text{Hom}(E^{\times \ell}, F)), \ell = 0, \dots, m\}.$$

We can similarly define the space of C_{fb}^k sections of a bundle.

With this in mind, we present the following strengthening of Definition 4.5.

Definition 4.9. *Fix $k \geq 2$. A k -tame perturbation q is called very k -tame if the following additional conditions are satisfied:*

- (i) \hat{q} defines an element of $C_{\text{fb}}^\infty(\mathcal{C}_k(Z), \mathcal{V}_k(Z))$;
 - (ii) \hat{q} also defines an element of $C_{\text{fb}}^0(\mathcal{C}_j(Z), \mathcal{V}_j(Z))$, for all integers $j \in [1, k]$;
 - (iii) $\mathcal{D}\hat{q}$ defines an element of $C_{\text{fb}}^\infty(\mathcal{C}_k(Z), \text{Hom}(T\mathcal{C}_j(Z), \mathcal{V}_j(Z)))$, for all integers $j \in [-k, k]$.
- We say that q is very tame if it is very k -tame for all $k \geq 2$.

Luckily, the kind of perturbations considered by Kronheimer and Mrowka in [KM07] satisfy these additional conditions.

Lemma 4.10. *Let q be the gradient of a cylinder function, as in [KM07, Section 11]. Then, q is very tame.*

Proof. This is an immediate consequence of Proposition 11.4.1 and Lemma 11.4.4 in [KM07], along the lines of the proof of Theorem 11.1.2 on p.190 of [KM07]. \square

We will need to ensure that both the perturbations we use are very tame and the space of perturbations is sufficiently large to guarantee transversality. This motivates the following definition.

Definition 4.11. *A large Banach space of very tame perturbations is a large Banach space of tame perturbations where each perturbation is additionally very tame.*

Using Lemma 4.10, the proof of the existence of such a space follows as in the proof of Theorem 11.6.1 in [KM07].

4.5. Admissible perturbations. For Morse homology, we need non-degeneracy of the Hessian at the critical points and the Morse-Smale condition. For monopole Floer homology, we must establish the analogues of these: non-degeneracy of stationary points and regularity of the moduli spaces. We first focus on non-degeneracy.

Fix a tame perturbation q and Sobolev number k . Note that \mathcal{X}_q^σ takes $\mathcal{C}_k^\sigma(Y)$ to \mathcal{T}_{k-1}^σ . Note also that the stationary points of \mathcal{X}_q^σ cannot be isolated like in Morse theory. This is because the gauge group preserves stationary points. Therefore, this will be a Morse-Bott condition, in the sense that the stationary points will be non-degenerate in the directions transverse to the gauge orbits.

Definition 4.12. *A stationary point x of \mathcal{X}_q^σ is non-degenerate in $\mathcal{C}_k^\sigma(Y)$ if \mathcal{X}_q^σ is transverse to the sub-bundle \mathcal{J}_{k-1}^σ at x .*

Much like the existence of stationary points, the non-degeneracy of the stationary points of \mathcal{X}_q^σ can be described in terms of \mathcal{X}_q on $\mathcal{C}_k(Y)$.

Proposition 4.13 (Proposition 12.2.5 in [KM07]). *A stationary point (a, s, ϕ) is a non-degenerate zero of \mathcal{X}_q^σ on $\mathcal{C}_k^\sigma(Y)$ for $s > 0$ if \mathcal{X}_q is transverse to \mathcal{J}_{k-1} at $(a, s\phi)$. If $s = 0$, and ϕ is an eigenvector of $D_{q,a}$ with eigenvalue λ , then $(a, 0, \phi)$ is non-degenerate if and only if the following three conditions are satisfied:*

- (1) $\lambda \neq 0$,
- (2) λ is a simple eigenvalue,
- (3) \mathcal{X}_q is transverse to $\mathcal{J}_{k-1}^{\text{red}}$ at a .

It turns out that one can always find many perturbations such that the stationary points are non-degenerate.

Theorem 4.14 (Theorem 12.1.2 in [KM07]). *For any large Banach space of tame perturbations \mathcal{P} , there is a residual subset of \mathcal{P} such that for each perturbation \mathfrak{q} , all of the stationary points of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ are non-degenerate.*

Given such a perturbation \mathfrak{q} , the gauge-equivalence classes of solutions to the perturbed Seiberg-Witten equations on (Y, \mathfrak{s}) are isolated (this comes from being transverse to the gauge orbits). Since these equivalence classes form a compact subset of $\mathcal{C}_k(Y)/\mathcal{G}_{k+1}$ because of tameness, there are at most finitely many irreducible solutions in $\mathcal{B}_k^{\sigma}(Y)$.

For the rest of the subsection, we will assume that \mathfrak{q} has been chosen so as to be tame and that the stationary points are non-degenerate. In analogy with the finite-dimensional case we have

Definition 4.15. *Let $[x] = [(a, s, \phi)]$ be a gauge-equivalence class of zeros of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ in $\mathcal{B}_k^{\sigma}(Y)$. If $s \neq 0$, then we say that $[x]$ is irreducible. If $[x]$ is reducible, then $[x]$ is boundary-stable (respectively boundary-unstable) if $\Lambda_{\mathfrak{q}}(x) > 0$ (respectively < 0).*

Define \mathfrak{C}^o , \mathfrak{C}^s , and \mathfrak{C}^u to be the sets of gauge-equivalence classes of irreducible, boundary-stable, and boundary-unstable stationary points of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ on $\mathcal{B}_k^{\sigma}(Y)$. Set $\mathfrak{C} = \mathfrak{C}^o \cup \mathfrak{C}^s \cup \mathfrak{C}^u$.

Next, we discuss regularity for trajectories. Let $Z = \mathbb{R} \times Y$. For $x, y \in \mathcal{C}^{\sigma}(Y)$, pick a smooth path $\gamma_0 : \mathbb{R} \rightarrow \mathcal{C}^{\sigma}(Y)$ such that $\gamma_0(t) = x$ for $t \ll 0$ and $\gamma_0(t) = y$ for $t \gg 0$. Then, define

$$(48) \quad \mathcal{C}_k^{\tau}(x, y) = \{\gamma \in \mathcal{C}_{k,loc}^{\tau}(Z) \mid \gamma - \gamma_0 \in L_k^2(Z; iT^*Z) \times L_k^2(\mathbb{R}; \mathbb{R}) \oplus L_k^2(Z; \mathbb{S}^+)\},$$

Note that here we are imposing L_k^2 and not $L_{k,loc}^2$ conditions on $\gamma - \gamma_0$. Thus, $\mathcal{C}_k^{\tau}(x, y)$ is equipped with a natural metric, $d(\gamma, \gamma') = \|\gamma - \gamma'\|_{L_k^2}$. We have an analogously defined space $\tilde{\mathcal{C}}_k^{\tau}(x, y)$ where we remove the condition $s(t) \geq 0$, as for the definition of $\tilde{\mathcal{C}}_k^{\tau}(Z)$.

We let $\mathcal{B}_k^{\tau}(x, y)$ (respectively $\tilde{\mathcal{B}}_k^{\tau}(x, y)$) be the quotient of $\mathcal{C}_k^{\tau}(x, y)$ (respectively $\tilde{\mathcal{C}}_k^{\tau}(x, y)$) by the action of gauge transformations $u : Z \rightarrow S^1$ such that $1 - u \in L_{k+1}^2(Z, \mathbb{C})$. The space $\mathcal{B}_k^{\tau}(x, y)$ depends only on the classes $[x]$ and $[y]$, up to canonical diffeomorphism. Consequently, we can use the notation $\mathcal{B}_k^{\tau}([x], [y])$. Similarly for $\tilde{\mathcal{B}}_k^{\tau}([x], [y])$.

We are now interested in studying gauge-equivalence classes of trajectories between two stationary points. Define $\tau_t : \mathbb{R} \times Y \rightarrow \mathbb{R} \times Y$ to be translation by t and let γ_x (respectively γ_y) denote the elements of $\mathcal{C}_{k,loc}^{\tau}(\mathbb{R} \times Y)$ which are x (respectively y) in each slice.

Definition 4.16. *For $[x]$ and $[y]$ in \mathfrak{C} , we define the moduli space of trajectories from $[x]$ to $[y]$ as*

$$M([x], [y]) = \{[\gamma] \in \mathcal{B}_{k,loc}^{\tau}(\mathbb{R} \times Y) \mid \mathcal{F}_{\mathfrak{q}}^{\tau}(\gamma) = 0, \lim_{t \rightarrow -\infty} [\tau_t^* \gamma] = [\gamma_x], \lim_{t \rightarrow +\infty} [\tau_t^* \gamma] = [\gamma_y]\},$$

where the limits are taken in $\mathcal{B}_{k,loc}^{\tau}(\mathbb{R} \times Y)$. If $[x]$ is boundary-stable and $[y]$ is boundary-unstable, then we say we are in the boundary-obstructed case. Finally, we will decorate a moduli space by \tilde{M} if we want the result after quotienting by the usual \mathbb{R} -action.

It is proved in [KM07, Theorem 13.3.5] that every class in $M([x], [y])$ has a gauge representative in $\mathcal{B}_k^{\tau}([x], [y])$. The proof uses exponential decay estimates for Seiberg-Witten trajectories.

Remark 4.17. As mentioned above, each (blown-up, perturbed) Seiberg-Witten trajectory on $\mathcal{C}_k^\sigma(Y)$ is gauge-equivalent to a smooth trajectory on $\mathcal{C}^\sigma(Y)$, for $k \geq 2$. In particular, the moduli spaces are homeomorphic for any $k \geq 2$.

In order to show that the moduli spaces are smooth manifolds, we must show that they are (locally) the preimage of a regular value of a Fredholm map. Just as in the case of compact cylinders, one can define a bundle $\mathcal{V}_{k-1}^\tau(Z)$ over $\tilde{\mathcal{C}}_k^\tau(x, y)$, such that \mathcal{F}_q^τ provides a section. The zero set of this section, restricted to $\mathcal{C}_k^\tau(x, y)$, thus describes the Seiberg-Witten trajectories asymptotic to x and y . The advantage of working with $\tilde{\mathcal{C}}_k^\tau(x, y)$ is that we can differentiate \mathcal{F}_q^τ in this setting.

Recall from (28) that $\mu_Y(f)$ denotes the average value of a function f over the three-manifold Y . For $\gamma = (a, s, \phi) \in \tilde{\mathcal{C}}_k^\tau(x, y)$ and $1 \leq j \leq k$, define the linear operator⁸

$$(49) \quad Q_\gamma = \mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau \oplus \mathbf{d}_\gamma^{\tau, \dagger} : \mathcal{T}_{j, \gamma}^\tau(Z) \rightarrow \mathcal{V}_{j-1, \gamma}^\tau(Z) \oplus L_{j-1}^2(Z; i\mathbb{R}),$$

where

$$(50) \quad \mathbf{d}_\gamma^{\tau, \dagger}(b, r, \psi) = -d^*b + is^2 \operatorname{Re}\langle i\phi, \psi \rangle + i|\phi|^2 \operatorname{Re} \mu_Y \langle i\phi, \psi \rangle$$

is a variant of the formal adjoint to the infinitesimal (four-dimensional) gauge action. The condition $\mathbf{d}_\gamma^{\tau, \dagger} = 0$ describes a four-dimensional local Coulomb slice $\mathcal{K}_{j, \gamma}^\tau \subset \mathcal{T}_{j, \gamma}^\tau$. When $j = k$, the slice $\mathcal{K}_{k, \gamma}^\tau$ can be viewed as the tangent space to $\tilde{\mathcal{B}}_k^\tau([x], [y])$ at $[\gamma]$. In general, the bundle \mathcal{K}_j^τ defines a gauge-equivariant bundle over $\tilde{\mathcal{C}}_k^\tau(x, y)$ which descends to a bundle over $\tilde{\mathcal{B}}_k^\tau([x], [y])$. If Q_γ is surjective (or, equivalently, the restriction $\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau|_{\mathcal{K}_{j, \gamma}^\tau}$ surjects onto $\mathcal{V}_{j-1, \gamma}^\tau$) for all $[\gamma]$ in $M([x], [y])$, then $M([x], [y])$ is a smooth manifold.

Definition 4.18. *The moduli space $M([x], [y])$ is regular if Q_γ is surjective for all $[\gamma] \in M([x], [y])$, unless we are in the boundary-obstructed case. If $M([x], [y])$ is boundary-obstructed, regularity means that the cokernel of Q_γ is dimension one for each $[\gamma] \in M([x], [y])$.*

Theorem 15.1.1 in [KM07] guarantees the existence of a tame perturbation \mathbf{q} such that $M([x], [y])$ are regular for all $[x]$ and $[y]$. In particular, this implies that the moduli spaces $M([x], [y])$ are all smooth manifolds, even in the boundary-obstructed case, by [KM07, Proposition 14.5.7]. In particular, we have that $\operatorname{ind} Q_\gamma = \dim M([x], [y])$, except in the boundary-obstructed case, where $\operatorname{ind} Q_\gamma = \dim M([x], [y]) - 1$.

If $[x]$ and $[y]$ are reducible and $M([x], [y])$ is regular, then define $M^{\operatorname{red}}([x], [y])$ to be the subset of $M([x], [y])$ consisting of reducible trajectories. Note that $M^{\operatorname{red}}([x], [y])$ is either empty (this is the case when one of $[x]$ or $[y]$ is irreducible) or all of $M([x], [y])$, except when $[x]$ is boundary-unstable and $[y]$ is boundary-stable, in which case it is $\partial M([x], [y])$.

Definition 4.19. *An admissible perturbation is a tame perturbation such that all the stationary points are non-degenerate and all the moduli spaces are regular.*

Theorem 4.20 (Theorem 15.1.1 in [KM07]). *For any large Banach space of tame perturbations \mathcal{P} , there exists an admissible perturbation $\mathbf{q} \in \mathcal{P}$.*

As discussed after Definition 4.11, there exist large Banach spaces of very tame perturbations. Therefore, there exist perturbations which are both very tame and admissible.

⁸In Equation (49), \mathcal{D}^τ is a covariant derivative on the space of paths, which was simply denoted \mathcal{D} in [KM07]. See Section 5.9 below for more details.

4.6. Orientations. From now on we fix an admissible perturbation \mathbf{q} .

Theorem 4.21 (Corollary 20.4.1 of [KM07]). *The moduli spaces $M([x], [y])$ are orientable manifolds.*

Let us sketch the construction of orientations on $M([x], [y])$, following [KM07, Section 20]. This is similar to the discussion of specialized coherent orientations in Morse theory (see Section 2.2).

To orient $M([x], [y])$, we need to orient the determinant lines $\det(Q_\gamma)$, where Q_γ is as in (49). For arbitrary $x, y \in \mathcal{C}_k^\sigma(Y)$ (not necessarily stationary points), we consider instead a compact interval $I = [t_1, t_2]$, and a configuration $\gamma \in \mathcal{C}_k^\tau(I \times Y)$ whose restrictions to $\{t_1\} \times Y$ and $\{t_2\} \times Y$ are gauge equivalent to x , resp. y . To any such γ we can associate an operator Q_γ by the same formula as for (49). To make it Fredholm, we need to add suitable boundary conditions. At the boundary component $\{t_1\} \times Y$, consider the subspaces

$$H_1^\pm = \mathcal{K}_{1/2,x}^\pm \oplus L_{1/2}^2(Y; i\mathbb{R}) \subset \mathcal{T}_{1/2,x}^\sigma(Y) \oplus L_{1/2}^2(Y; i\mathbb{R}).$$

Here, we have decomposed the local Coulomb slice $\mathcal{K}_{1/2,x}^\sigma \subset \mathcal{T}_{1/2,x}^\sigma(Y)$ as

$$\mathcal{K}_{1/2,x}^\sigma = \mathcal{K}_{1/2,x}^- \oplus \mathcal{K}_{1/2,x}^+,$$

using the spectral subspaces (the direct sum of nonpositive, resp. positive eigenspaces) of the Hessian

$$\text{Hess}_{\mathbf{q},x}^\sigma := \Pi_{\mathcal{K}_{1/2,x}^\sigma} \circ \mathcal{D}_\gamma \mathcal{X}_{\mathbf{q}}^\sigma|_{\mathcal{K}_{1/2,x}^\sigma}.$$

(See Section 5.5.2 below for more details about the Hessian.)

We consider similar subspaces \mathcal{K}_2^\pm at the other boundary component $\{t_2\} \times Y$. Then, we define a Fredholm operator

$$(51) \quad P_\gamma = (Q_\gamma, -\Pi_1^+, \Pi_2^-) : \mathcal{T}_{1,\gamma}^\tau(I \times Y) \rightarrow \mathcal{V}_{0,\gamma}^\tau \oplus L^2(I \times Y; i\mathbb{R}) \oplus H_1^+ \oplus H_2^-,$$

where Π_i^\pm denotes the composition of restriction to $\{t_i\} \times Y$ with the projection to H_i^\pm .

Let $\Lambda(\gamma)$ be the set of orientations of $\det(P_\gamma)$. It is proved in [KM07, Proposition 20.3.4] that there are canonical identifications between the different $\Lambda(\gamma)$ when we fix the (gauge equivalence classes of the) endpoints x and y of γ . Thus, we can write $\Lambda([x], [y])$ for $\Lambda(\gamma)$. Further, there are natural composition maps

$$\Lambda([x], [y]) \times \Lambda([y], [z]) \rightarrow \Lambda([x], [z]).$$

Departing slightly from the terminology and conventions in [KM07], we define an *orientation data set* o for the admissible perturbation \mathbf{q} to consist of elements $o_{[x],[y]} \in \Lambda([x], [y])$, one for each pair $([x], [y])$, such that we have the relations

$$o_{[x],[y]} \cdot o_{[y],[z]} = o_{[x],[z]}.$$

An orientation data set can be constructed as follows. In [KM07, p. 385-390] it is shown that when $[x]$ and $[y]$ are reducible, the set $\Lambda([x], [y])$ can be canonically identified with $\mathbb{Z}/2 = \{\pm 1\}$, in a way compatible with concatenation; we then choose $o_{[x],[y]}$ to be the element $+1$ in this case. Then, we fix a reducible $[x_0]$ and pick arbitrary elements $o_{[x_0],[x]}$ for all irreducibles $[x]$. This uniquely determines the data set o , using the concatenation property.

Next, let $[x]$ and $[y]$ be stationary points, and consider a trajectory $[\gamma] \in M([x], [y])$. Let γ_0 be its restriction to a large compact interval $I = [t_1, t_2] \subset \mathbb{R}$. It is proved in [KM07, Section 20.4] that an orientation for $\det(P_{[\gamma_0]})$ determines one for $\det(Q_\gamma)$. Therefore, an

orientation data set for \mathbf{q} produces orientations for all determinant lines $\det(Q_\gamma)$, and hence for the moduli spaces $M([x], [y])$. From here we also get orientations on the quotients $\check{M}([x], [y])$.

We can orient the moduli spaces $M^{\text{red}}([x], [y])$ and $\check{M}^{\text{red}}([x], [y])$ in a similar manner.

Remark 4.22. In [KM07, Section 20], the discussion was more general, allowing the perturbation \mathbf{q} to vary for $t \in [t_1, t_2]$. For the purposes of this paper, it suffices to work with a single perturbation \mathbf{q} ; we adjusted our discussion accordingly, for simplicity.

Furthermore, the notion of orientation data set did not appear in [KM07]. Whereas we trivialize the sets $\Lambda([x], [y])$ and then let the monopole Floer complex be generated by stationary points $[x]$ (cf. Section 4.7 below), Kronheimer and Mrowka define the complex more canonically, as generated by orientations of $\Lambda([x_0], [x])$ (where $[x_0]$ is a fixed reducible basepoint), with two opposite orientations being set to be the negative of each other. The two definitions are readily seen to be equivalent.

4.7. Monopole Floer homology. We will now define the monopole Floer chain complex $\widetilde{CM}(Y, \mathfrak{s}, \mathbf{q})$ analogous to the construction of Morse homology for manifolds with boundary in Section 2.5.

The set \mathfrak{C}^o of gauge-equivalence classes of irreducible stationary points of $\mathcal{X}_\mathbf{q}^\sigma$ is finite. There are also countably many reducible stationary points, corresponding to the eigenvectors of $D_{\mathbf{q},a}$. Indeed, the operator $D_{\mathbf{q},a}$ is ASAFOE in the sense of [KM07, Definition 12.2.1], so its spectrum is discrete by [KM07, Lemma 12.2.4].

In [KM07, Chapter V], Kronheimer and Mrowka give an analysis of the compactifications of the moduli spaces via broken flow lines. Rather than state the general result, we simply point out that given an admissible perturbation, if $\check{M}([x], [y])$ or $\check{M}^{\text{red}}([x], [y])$ is 0-dimensional, then this moduli space is compact.

Let the groups C^θ be freely generated over \mathbb{Z} by \mathfrak{C}^θ for $\theta \in \{o, s, u\}$. The monopole Floer chain groups are given by

$$\widetilde{CM}(Y, \mathfrak{s}, \mathbf{q}) = C^o \oplus C^s.$$

Here, the relative \mathbb{Z} -grading on $\widetilde{CM}(Y, \mathfrak{s}, \mathbf{q})$ (since \mathfrak{s} is torsion) is given by the expected dimension of $M([x], [y])$ (the index of Q_γ), since in this case, we cannot have that $M([x], [y])$ is boundary-obstructed. Since Y is a rational homology sphere, this number is well-defined (there is only one homotopy class of paths from $[x]$ to $[y]$). Recall that we need a shift of grading by 1 in the boundary-obstructed case.

For $\theta, \varpi \in \{o, s, u\}$, define ∂_ϖ^θ and $\bar{\partial}_\varpi^\theta$ by

$$(52) \quad \partial_\varpi^\theta([x]) = \sum_{[y] \in \mathfrak{C}^\varpi} \# \check{M}([x], [y])[y],$$

$$(53) \quad \bar{\partial}_\varpi^\theta([x]) = \sum_{[y] \in \mathfrak{C}^\varpi} \# \check{M}^{\text{red}}([x], [y])[y],$$

for $[x] \in \mathfrak{C}^\theta$, where we only sum over $[y]$ such that the relevant moduli spaces are 0-dimensional. Here, $\#$ means the signed count of points in this oriented, compact 0-dimensional manifold. Also, we only allow θ, ϖ such that these counts make sense: for ∂ , we want $\theta \in \{o, u\}$ and $\varpi \in \{o, s\}$ while for $\bar{\partial}$, we ask that $\theta, \varpi \neq o$. The boundary operator on $\widetilde{CM}(Y, \mathfrak{s}, \mathbf{q})$ is given by

$$(54) \quad \check{\partial} = \begin{bmatrix} \partial_o^o & -\partial_o^u \bar{\partial}_u^s \\ \partial_s^o & \bar{\partial}_s^s - \partial_s^u \bar{\partial}_u^s \end{bmatrix}.$$

The admissibility of \mathfrak{q} guarantees that $\tilde{\partial}$ squares to zero and we can take homology, $\widetilde{HM}(Y, \mathfrak{s}, \mathfrak{q})$. As in Remark 2.17, the moduli spaces $M^{\text{red}}([x], [y])$ are either empty or equal to $M([x], [y])$, except in the case that $[x]$ is boundary-unstable and $[y]$ is boundary-stable; in this exceptional case, the counts of $\check{M}^{\text{red}}([x], [y])$ do not arise in $\tilde{\partial}$.

Similar to the Morse homology for circle actions, monopole Floer homology can be given the structure of a $\mathbb{Z}[U]$ -module. This is defined in [KM07] in terms of evaluations of suitable Čech cochains on the 2-dimensional moduli spaces of parameterized trajectories. For our purposes, it is more convenient to use an alternate, equivalent definition of the U map, which is taken from [KMOS07, Section 4.11]. The latter definition involves counting points in zero-dimensional spaces, and this will make it easier to prove a stability result for trajectories in Section 13.2.

Although the U map on \widetilde{HM} comes from a more general cobordism construction, we will restrict our attention to the specific case we need. Let $p \in \mathbb{R} \times Y$ be a basepoint, and B_p a standard ball neighborhood of p in $\mathbb{R} \times Y$. Let $\mathcal{B}_k^\sigma(B_p)$ be the blown-up configuration space of L_k^2 connections and spinors on B_p , modulo the gauge group

$$\mathcal{G}_{k+1}(B_p) = \{u : B_p \rightarrow S^1 \mid u \in L_{k+1}^2\}.$$

Note that $\mathcal{B}_k^\sigma(B_p)$ is a Hilbert manifold with boundary, and is a free quotient by $\mathcal{G}_{k+1}(B_p)$. There is a natural complex line bundle E_p^σ over $\mathcal{B}_k^\sigma(B_p)$, induced from the map $\mathcal{G}_{k+1}(B_p) \rightarrow S^1$, $u \mapsto u(p)$.

For any $[x] \in \mathfrak{C}^\theta$, $[y] \in \mathfrak{C}^\varpi$, because of unique continuation, there is a well-defined restriction map

$$r_p : M([x], [y]) \rightarrow \mathcal{B}_k^\sigma(B_p),$$

which is an embedding. Pick a smooth section ζ of E_p^σ such that ζ is transverse to the zero section, and the zero set \mathcal{Z} of ζ intersects all the moduli spaces $M([x], [y])$ and $M^{\text{red}}([x], [y])$ transversely. By analogy with the finite-dimensional case from Section 2.7, for $(\theta, \varpi) \in \{(o, o), (o, s), (u, o), (u, s)\}$, we define $m_\varpi^\theta : \mathfrak{C}^\theta \rightarrow \mathfrak{C}^\varpi$ by

$$(55) \quad m_\varpi^\theta([x]) = \sum_{[y] \in \mathfrak{C}^\varpi} \#(M([x], [y]) \cap \mathcal{Z}) \cdot [y], \text{ for } [x] \in \mathfrak{C}^\theta,$$

Similarly, for the reducibles, we set

$$(56) \quad \bar{m}_\varpi^\theta([x]) = \sum_{[y] \in \mathfrak{C}^\varpi} \#(M^{\text{red}}([x], [y]) \cap \partial \mathcal{Z}) \cdot [y], \text{ for } [x] \in \mathfrak{C}^\theta.$$

Note that we are using parameterized trajectories. We only consider terms in the sums where the dimension of the moduli spaces being considered is two.

Finally, we can define $\check{m} : \widetilde{CM} \rightarrow \widetilde{CM}$ by

$$(57) \quad \check{m} = \begin{bmatrix} m_o^o & -m_o^u \bar{\partial}_u^s - \partial_o^u \bar{m}_u^s \\ m_s^o & \bar{m}_s^s - m_s^u \bar{\partial}_u^s - \partial_s^u \bar{m}_u^s \end{bmatrix}.$$

This induces a chain map. Therefore, we can define the U -action on \widetilde{HM} by the map induced by \check{m} . By construction, this map lowers the relative grading by 2.

Theorem 4.23 (Kronheimer-Mrowka, Theorem 23.1.5 of [KM07]). *The relatively-graded $\mathbb{Z}[U]$ -module $\widetilde{HM}(Y, \mathfrak{s}, \mathfrak{q})$ is an invariant of the pair (Y, \mathfrak{s}) .*

Therefore, we just use the notation $\widetilde{HM}(Y, \mathfrak{s})$. From here we can easily define Bloom's variant of monopole Floer homology, \widehat{HM} .

Definition 4.24 ([Blo11]). *The tilde-flavor of monopole Floer homology, $\widetilde{HM}(Y, \mathfrak{s})$, is the homology of the mapping cone of the map \widetilde{m} on $\widetilde{CM}(Y, \mathfrak{s}, \mathfrak{q})$.*

4.8. Gradings. Recall that for two stationary points x and y , the relative homological grading between $[x]$ and $[y]$ is defined by

$$\text{gr}([x], [y]) = \text{ind } Q_\gamma,$$

where $\gamma \in \mathcal{C}_k^T(x, y)$ is any path, and Q_γ is the operator from (49). (See [KM07, Definition 14.4.4].)

Suppose that $[x]$ and $[y]$ are reducibles with the same blow-down projection $(a, 0)$, and correspond to eigenvalues μ and ν of $D_{\mathfrak{q}, a}$ with $\mu \geq \nu$. Then, an alternate formula for the relative grading is given in [KM07, Corollary 14.6.2]:

$$(58) \quad \text{gr}([x], [y]) = \begin{cases} 2i(\mu, \nu) & \text{if } \mu \text{ and } \nu \text{ have the same sign,} \\ 2i(\mu, \nu) - 1 & \text{otherwise.} \end{cases}$$

where $i(\mu, \nu)$ denotes the number of eigenvalues in the interval $(\nu, \mu]$.

There is also an absolute rational grading $\text{gr}^\mathbb{Q}$ defined in [KM07, Section 28.3]. Given a cobordism W from the round sphere S^3 to Y , we have

$$\text{gr}^\mathbb{Q}([x]) := -\text{gr}_z([x_0], W, [x]) + \frac{c_1(\mathfrak{t})^2 - \sigma(W)}{4} - \iota(W),$$

where: $[x_0]$ is the reducible on S^3 with the lowest positive eigenvalue; \mathfrak{t} is a Spin^c structure on W that restricts to \mathfrak{s} on Y ; $\text{gr}_z([x_0], W, [x])$ is the relative grading between x_0 and x for the monopole map associated to (W, \mathfrak{t}) , as in [KM07, Section 25]; the subscript z refers to a connected component of the configuration space of W , and in our case z is uniquely determined by the Spin^c structure \mathfrak{t} ; and $\iota(W) = (\chi(W) + \sigma(W))/2 = b^+(W) - b_1(W)$.

For future reference, let us rephrase this definition in terms of a four-manifold X with boundary Y . We can then take W to be the complement of a four-ball in X . Using additivity of gr and a standard calculation on B^4 we can replace W with X and write

$$(59) \quad \text{gr}^\mathbb{Q}([x]) = -\text{gr}_z(X, [x]) + \frac{c_1(\mathfrak{t})^2 - \sigma(X)}{4} - b^+(X) + b_1(X) - 1,$$

where $\text{gr}_z(X, [x])$ is the expected dimension of the moduli space $M_z(X^*; [x])$ defined in [KM07, Section 24]. Here, X^* refers to X after attaching to it a cylindrical end of the form $[0, \infty) \times Y$.

5. REDUCTION TO THE COULOMB GAUGE

In this section, we recast the constructions from Section 4 entirely in the global Coulomb slice

$$W = \ker d^* \oplus \Gamma(\mathbb{S}) \subset \mathcal{C}(Y).$$

This is needed in order to make contact with the construction of the Seiberg-Witten Floer spectrum from Section 3, for which we used finite dimensional approximation on W . Throughout this section k will denote a fixed integer at least 2.

5.1. Bundle decompositions and projections. Recall that in Section 3.2 we introduced the global Coulomb slice W , the global Coulomb projection Π^{gC} , the infinitesimal global Coulomb projection Π_*^{gC} , the local Coulomb slice \mathcal{K} , the (infinitesimal) local Coulomb projection Π^{lC} , the enlarged local Coulomb slice \mathcal{K}^e , the enlarged local Coulomb projection Π^{elC} , and the metric \tilde{g} on W . Further, in Section 4.1 we introduced the bundle decomposition of \mathcal{T}_k into the tangents to the gauge orbits \mathcal{J}_k and the local Coulomb slice \mathcal{K}_k ; we also mentioned a similar decomposition in the blow-up, $\mathcal{T}_k^\sigma = \mathcal{J}_k^\sigma \oplus \mathcal{K}_k^\sigma$.

In this section we explore these constructions further. In particular, we extend the gauge projections to the blow-up, and describe a few bundle decompositions that are related to global Coulomb gauge.

For $j \leq k$, recall that $\mathcal{T}_j^{\text{gC}}$ is the trivial vector bundle with fiber W_j over W_k . Unlike what we did for \mathcal{T}_j in Section 4.3, in the case of $\mathcal{T}_j^{\text{gC}}$ there is no need to define the Sobolev norms using the varying covariant derivatives ∇_{A_0+a} . The Sobolev norm on W_j defined by $\nabla = \nabla_{A_0}$ is invariant under the residual gauge action by S^1 , and this norm is what we shall use on each tangent space. Thus, $\mathcal{T}_j^{\text{gC}}$ is exactly the trivial normed bundle $W_k \times W_j$. Again, we keep the notation $\mathcal{T}_j^{\text{gC}}$ (rather than just $W_k \times W_j$) to emphasize the bundle structure. Note that these two Sobolev norms are equivalent in the following strong sense. For $x = (a, \phi) \in W_k$ with $\|a\|_{L_k^2} < R$, there exists a constant $C(R)$ such that

$$\frac{1}{C(R)} \|v\|_{L_j^2} \leq \|v\|_{L_{j,a}^2} \leq C(R) \|v\|_{L_j^2}$$

for all $v \in \mathcal{T}_{j,x}$ and $j \leq k$. Therefore, these Sobolev norms will also be equivalent when working over paths of three-dimensional configurations, and so even in the four-dimensional setting, we are content to use ∇ .

For $x = (a, \phi) \in W_k$, we let $\mathcal{J}_x^{\text{gC}} \subset \mathcal{T}_k^{\text{gC}}$ be the (real) span of $(0, i\phi)$. This is the tangent space to the S^1 -orbit at x . Since $\mathcal{J}_x^{\text{gC}}$ is canonically isomorphic to \mathbb{R} , there is no need for a subscript j for Sobolev regularity.

The intersection of a gauge orbit with global Coulomb gauge is depicted schematically in Figure 1.

Lemma 5.1. *Fix $j \geq 0$. For each $x = (a, \phi) \in W_k$ with $\phi \neq 0$, we have $\mathcal{J}_{j,x} \cap \mathcal{T}_{j,x}^{\text{gC}} = (\Pi_*^{\text{gC}})_x(\mathcal{J}_{j,x}) = \mathcal{J}_x^{\text{gC}}$.*

Proof. Let $x = (a, \phi) \in W_{k-1}$. We recall from [KM07, Page 140] that

$$\mathcal{J}_{j,x} = \{(-d\zeta, \zeta\phi) \mid \zeta \in L_{j+1}^2(Y; i\mathbb{R})\}.$$

First, we compute $\mathcal{J}_{j,x} \cap \mathcal{T}_{j,x}^{\text{gC}}$. Suppose $(b, \psi) \in \mathcal{J}_{j,x} \cap \mathcal{T}_{j,x}^{\text{gC}}$. If we write $(b, \psi) = (-d\zeta, \zeta\phi)$, we have $d^*d\zeta = 0$. Since Y is a rational homology sphere, this implies that ζ is constant.




FIGURE 1. We show the gauge orbit $\mathcal{G}x$ intersecting W in a circle. Also shown are the tangent spaces to the orbits, and the global Coulomb projection Π^{gC} .

Therefore,

$$\mathcal{J}_{j,x} \cap \mathcal{T}_{j,x}^{\text{gC}} = \{(0, it\phi) \mid t \in \mathbb{R}\} = \mathcal{J}_x^{\text{gC}}.$$

Now, we study $(\Pi_*^{\text{gC}})_x(\mathcal{J}_{j,x})$. Let $(-d\zeta, \zeta\phi) \in \mathcal{J}_{j,x}$. Recall that

$$(\Pi_*^{\text{gC}})_x(-d\zeta, \zeta\phi) = (-d\zeta - d\xi, \zeta\phi + \xi\phi),$$

where $\xi : Y \rightarrow i\mathbb{R}$ satisfies $d^*(-d\zeta - d\xi) = 0$ and $\int_Y \xi = 0$. Again, since Y is a rational homology sphere, this implies that $\zeta + \xi$ is a constant it for some $t \in \mathbb{R}$. Therefore, we obtain

$$(\Pi_*^{\text{gC}})_x(-d\zeta, \zeta\phi) = (0, it\phi).$$

This implies that $(\Pi_*^{\text{gC}})_x(\mathcal{J}_{j,x}) = \mathcal{J}_x^{\text{gC}}$, as desired. \square

Analogous to the splitting of \mathcal{T}_j as $\mathcal{J}_j \oplus \mathcal{K}_j$, there exists a decomposition

$$(60) \quad \mathcal{T}_j^{\text{gC}} = \mathcal{J}_j^{\text{gC}} \oplus \mathcal{K}_j^{\text{agC}}.$$

Here, $\mathcal{K}_j^{\text{agC}}$ is defined to be the orthogonal complement of $\mathcal{J}_j^{\text{gC}}$ with respect to the \tilde{g} -metric on W . For each $x = (a, \phi) \in W_k$ and $j \leq k$, this space can be written explicitly as

$$\begin{aligned} \mathcal{K}_{j,x}^{\text{agC}} &= \{(b, \psi) \in \mathcal{T}_{j,x}^{\text{gC}} \mid \langle (0, i\phi), (b, \psi) \rangle_{\tilde{g}} = 0\} \\ &= \{(b, \psi) \in \mathcal{T}_{j,x} \mid d^*b = 0, \langle (0, i\phi), (b, \psi) \rangle_{\tilde{g}} = 0\}. \end{aligned}$$

We call $\mathcal{K}_{j,x}^{\text{agC}}$ the *anticircular global Coulomb slice* at a point $x = (a, \phi) \in W_k$. Observe that, if $d^*b = 0$, then the vector $(b, 0)$ is L^2 -perpendicular to all $(-d\zeta, \zeta\phi) \in \mathcal{J}_{j,x}$, and hence lies in the local Coulomb slice $\mathcal{K}_{j,x} \subset \mathcal{K}_{j,x}^e$. Using the formula (37) for the \tilde{g} -inner product, we get that

$$(61) \quad \langle (0, i\phi), (b, 0) \rangle_{\tilde{g}} = \text{Re} \langle (0, i\phi), \Pi^{\text{elC}}(b, 0) \rangle_{L^2} = \text{Re} \langle (0, i\phi), (b, 0) \rangle_{L^2} = 0.$$

We deduce that the condition $\langle (0, i\phi), (b, \psi) \rangle_{\tilde{g}} = 0$ is equivalent to

$$\langle (0, i\phi), (0, \psi) \rangle_{\tilde{g}} = 0.$$




FIGURE 2. Different Coulomb slices and projections inside the tangent space $\mathcal{T}_{j,x}$. We drop the subscripts j and x from notation for convenience.

For simplicity, we will write $\langle \psi_1, \psi_2 \rangle_{\tilde{g}}$ and $\|\psi\|_{\tilde{g}}^2$ for expressions of the form $\langle (0, \psi_1), (0, \psi_2) \rangle_{\tilde{g}}$ and $\|(0, \psi)\|_{\tilde{g}}^2$. With this in mind, we can write the anticircular global Coulomb slice as

$$\begin{aligned} \mathcal{K}_{j,x}^{\text{agC}} &= \{(b, \psi) \in \mathcal{T}_{j,x}^{\text{gC}} \mid \langle i\phi, \psi \rangle_{\tilde{g}} = 0\} \\ &= \{(b, \psi) \in \mathcal{T}_{j,x} \mid d^*b = 0, \langle i\phi, \psi \rangle_{\tilde{g}} = 0\}. \end{aligned}$$

We define

$$\Pi^{\text{agC}} = \Pi_{\mathcal{K}_j^{\text{agC}}} \circ \Pi_*^{\text{gC}} : \mathcal{T}_j \rightarrow \mathcal{K}_j^{\text{agC}}$$

to be the composition of infinitesimal global Coulomb projection $\Pi_*^{\text{gC}} : \mathcal{T}_j \rightarrow \mathcal{T}_j^{\text{gC}}$ with \tilde{g} orthogonal projection onto $\mathcal{K}_j^{\text{agC}}$. See Figure 2.

Remark 5.2. The anticircular global Coulomb slice $\mathcal{K}_{j,x}^{\text{agC}}$ contains, but does not equal, the intersection $\mathcal{K}_{j,x} \cap \mathcal{T}_j^{\text{gC}}$. Indeed, the latter consists of vectors (b, ψ) satisfying $d^*b = 0$ and $\text{Re}\langle i\phi, \psi \rangle = 0$ (pointwise). These conditions imply

$$\langle (0, i\phi), (b, \psi) \rangle_{\tilde{g}} = \text{Re}\langle (0, i\phi), \Pi_x^{\text{elC}}(b, \psi) \rangle_{L^2} = \text{Re}\langle (0, i\phi), (b, \psi) \rangle_{L^2} = 0,$$

and thus $(b, \psi) \in \mathcal{K}_{j,x}^{\text{agC}}$. Here we are using that Π^{elC} is an L^2 projection, which can easily be verified.

Remark 5.3. The subspaces \mathcal{J}^{gC} and $\mathcal{K}_j^{\text{agC}}$ do not form Hilbert bundles over the whole of W_k , because \mathcal{J}^{gC} is smaller (and $\mathcal{K}_j^{\text{agC}}$ larger) at reducibles, compared to irreducibles.

Recall that on W we have a natural metric \tilde{g} , defined by measuring the L^2 norm of the enlarged local Coulomb projections; cf. Equation (36). At this point it is helpful to extend the \tilde{g} -inner product on $\mathcal{T}_j^{\text{gC}}$ to the bigger bundle \mathcal{T}_j (restricted to W_k). Consider the bundle decomposition over W_k ,

$$(62) \quad \mathcal{T}_j = \mathcal{J}_j^{\circ} \oplus \mathcal{T}_j^{\text{gC}},$$

where \mathcal{J}_j° is the tangent to the orbit of the normalized gauge group \mathcal{G}° ; that is, $\mathcal{J}_j^\circ \subset \mathcal{J}_j$ consists of vectors $(-d\zeta, \zeta\phi)$ with $\int_Y \zeta = 0$. Given $x \in W_k$ and vectors $v, w \in \mathcal{T}_{j,x}$, decompose them according to (62) as

$$(63) \quad v = v^\circ + v^{\text{gC}}, \quad w = w^\circ + w^{\text{gC}},$$

where $v^{\text{gC}} = (\Pi_*^{\text{gC}})_x(v)$ and $w^{\text{gC}} = (\Pi_*^{\text{gC}})_x(w)$. Then, set

$$(64) \quad \langle v, w \rangle_{\tilde{g}} = \text{Re} \langle v^\circ, w^\circ \rangle_{L^2} + \langle v^{\text{gC}}, w^{\text{gC}} \rangle_{\tilde{g}},$$

where in the last inner product we use the formula (36). With this definition, the direct sum decompositions (62) and

$$(65) \quad \mathcal{T}_j = \mathcal{J}_j \oplus \mathcal{K}_j^{\text{agC}}$$

are orthogonal for \tilde{g} . Moreover, we can think of Π_*^{gC} and Π^{agC} as the \tilde{g} -orthogonal projections from \mathcal{T}_j to $\mathcal{T}_j^{\text{gC}}$ and $\mathcal{K}_j^{\text{agC}}$, respectively.

Next, we consider the blow-up $\mathcal{C}_k^\sigma(Y)$ from Section 4.1. This has tangent bundle \mathcal{T}^σ with Sobolev completions \mathcal{T}_j^σ . The tangents to the gauge orbits at $x = (a, s, \phi)$ form the subspace $\mathcal{J}_{x,j}^\sigma$, consisting of vectors of the form $(-d\zeta, 0, \zeta\phi)$. We also have a local Coulomb slice in the blow-up, given by the following formula from [KM07, Definition 9.3.6]:

$$\mathcal{K}_{j,x}^\sigma = \{(b, t, \psi) \in \mathcal{T}_{j,x}^\sigma \mid -d^*b + is^2 \text{Re} \langle i\phi, \psi \rangle = 0, \text{Re} \langle i\phi, \psi \rangle_{L^2} = 0\}.$$

The local Coulomb projection on the blow-up, $\Pi_{(a,s,\phi)}^{\text{IC},\sigma} : \mathcal{T}_{j,x}^\sigma \rightarrow \mathcal{K}_{j,x}^\sigma$, is defined to be

$$(66) \quad \Pi_{(a,s,\phi)}^{\text{IC},\sigma}(b, r, \psi) = (b - d\zeta, r, \psi + \zeta\phi),$$

where ζ is such that $-d^*(b - d\zeta) + is^2 \text{Re} \langle i\phi, \psi + \zeta\phi \rangle = 0$ and $\text{Re} \langle i\phi, \psi + \zeta\phi \rangle_{L^2} = 0$.

On the blow-up $\mathcal{C}^\sigma(Y)$ there is a natural L^2 metric, obtained from the inclusion $\mathcal{T}_j^\sigma \subset L_j^2(Y; iT^*Y) \oplus \mathbb{R} \oplus L_j^2(Y; \mathbb{S})$. Note that the direct sum decomposition

$$\mathcal{T}_j^\sigma = \mathcal{J}_j^\sigma \oplus \mathcal{K}_j^\sigma$$

is not orthogonal with respect to this L^2 metric. Rather, on the irreducible locus, the decomposition is orthogonal with respect to the pull-back of the L^2 metric on $\mathcal{C}(Y)$. On the other hand, this pull-back does not produce a non-degenerate metric on the whole $\mathcal{C}^\sigma(Y)$.

Now consider the residual gauge action of S^1 on W . It is convenient to blow up W at its fixed locus, just as we did with the configuration space $\mathcal{C}(Y)$. The blow-up of W is the space

$$W^\sigma = \{(a, s, \phi) \mid d^*a = 0, s \geq 0, \|\phi\|_{L^2} = 1\} \subset \mathcal{C}^\sigma(Y).$$

There is a natural identification of W^σ/S^1 with $\mathcal{B}^\sigma(Y)$, defined in Section 4.1. We have Sobolev completions W_k^σ and tangent bundles $\mathcal{T}_j^{\text{gC},\sigma}$ for $j \leq k$.

The global Coulomb projection $\Pi^{\text{gC}} : \mathcal{C}(Y) \rightarrow W$ induces a global Coulomb projection between the blow-ups:

$$(67) \quad \Pi^{\text{gC},\sigma} : \mathcal{C}^\sigma(Y) \rightarrow W^\sigma, \quad (a, s, \phi) \mapsto (a - df, s, e^f \phi),$$

where $f = Gd^*a$. Furthermore, if $x = (a, s, \phi) \in W^\sigma$, then the differential of $\Pi^{\text{gC},\sigma}$ is given by

$$(68) \quad (\Pi_*^{\text{gC},\sigma})_x : \mathcal{T}_x^\sigma \rightarrow \mathcal{T}_x^{\text{gC},\sigma}, \quad ((a, s, \phi), (b, r, \psi)) \mapsto (\pi(b), r, \psi + (Gd^*b)\phi).$$

At a point $x = (a, s, \phi) \in W^\sigma$, we let $\mathcal{J}_x^{\text{gC},\sigma} = \mathbb{R}\langle(0, 0, i\phi)\rangle$ be the tangent to the residual S^1 gauge orbit. Lemma 5.1 can be easily adapted to show that:

$$(69) \quad \mathcal{J}_{j,x}^\sigma \cap \mathcal{T}_{j,x}^{\text{gC},\sigma} = (\Pi_*^{\text{gC},\sigma})_x(\mathcal{J}_{j,x}^\sigma) = \mathcal{J}_x^{\text{gC},\sigma}.$$

Next, we define the anticircular global Coulomb slice in the blow-up. For $x = (a, s, \phi) \in W_k^\sigma$, let

$$(70) \quad \begin{aligned} \mathcal{K}_{j,x}^{\text{agC},\sigma} &= \{(b, r, \psi) \in \mathcal{T}_{j,x}^{\text{gC},\sigma} \mid \langle i\phi, \psi \rangle_{\tilde{g}} = 0\} \\ &= \{(b, r, \psi) \in \mathcal{T}_{j,x}^\sigma \mid d^*b = 0, \langle i\phi, \psi \rangle_{\tilde{g}} = 0\}, \end{aligned}$$

Note that the condition $(b, r, \psi) \in \mathcal{T}_{j,x}^\sigma$ already implies that $\text{Re}\langle\phi, \psi\rangle_{L^2} = 0$. Furthermore, here and later, by $\langle\psi, i\phi\rangle_{\tilde{g}}$ we implicitly mean that the inner product is taken in the blow-down projection; that is, we consider $\langle i\phi, \psi \rangle_{\tilde{g}(a,s,\phi)}$. It is useful to compare $\langle i\phi, \psi \rangle_{\tilde{g}}$ with the result of taking the \tilde{g} -inner product of $(0, i\phi)$ with the image of (b, r, ψ) in the blow-down: $(b, r\phi + s\psi)$. This yields

$$(71) \quad \langle (b, r\phi + s\psi), (0, i\phi) \rangle_{\tilde{g}} = s\langle (0, \psi), (0, i\phi) \rangle_{\tilde{g}} + r\langle (0, \phi), (0, i\phi) \rangle_{\tilde{g}} = s\langle \psi, i\phi \rangle_{\tilde{g}},$$

since $(0, \phi) \in \mathcal{K}_{(a,s,\phi)}$ and $\text{Re}\langle\phi, i\phi\rangle_{L^2} = 0$. Note that if we tried to define $\mathcal{K}_{(a,s,\phi)}^{\text{agC},\sigma}$ using the \tilde{g} -inner product of $(0, i\phi)$ with $(b, r\phi + s\psi)$, we would not obtain a bundle, since this would be larger at reducibles. On the other hand, for irreducibles, we have the following.

Lemma 5.4. *At every irreducible $x = (a, s, \phi) \in W_k^\sigma$, the infinitesimal blow-down projection*

$$(b, r, \psi) \mapsto (b, r\phi + s\psi)$$

induces a linear isomorphism from $\mathcal{K}_{j,(a,s,\phi)}^{\text{agC},\sigma}$ to $\mathcal{K}_{j,(a,s,\phi)}^{\text{agC}}$.

Proof. This follows from (71), since at an irreducible we have $s \neq 0$. \square

We have direct sum decompositions

$$(72) \quad \mathcal{T}_{j,x}^\sigma = \mathcal{J}_{j,x}^\sigma \oplus \mathcal{K}_{j,x}^{\text{agC},\sigma}$$

and

$$(73) \quad \mathcal{T}_{j,x}^{\text{gC},\sigma} = \mathcal{J}_x^{\text{gC},\sigma} \oplus \mathcal{K}_{j,x}^{\text{agC},\sigma}.$$

We define the anticircular global Coulomb projection on the blow-up,

$$\Pi_x^{\text{agC},\sigma} : \mathcal{T}_{j,x}^\sigma \rightarrow \mathcal{K}_{j,x}^{\text{agC},\sigma},$$

to be the projection with kernel $\mathcal{J}_{j,x}^\sigma$. It can be viewed as the composition of the map $(\Pi_*^{\text{gC},\sigma})_x$ from (68) with the projection with kernel $\mathcal{J}_x^{\text{gC},\sigma}$. This last projection, which is the restriction of $\Pi_x^{\text{agC},\sigma}$ to $\mathcal{T}_{j,x}^{\text{gC},\sigma}$, can be written explicitly as

$$(74) \quad (b, r, \psi) \mapsto (b, r, \psi) - \frac{\langle i\phi, \psi \rangle_{\tilde{g}}}{\|i\phi\|_{\tilde{g}}^2} \cdot (0, 0, i\phi).$$

From (72) we see that the anticircular global Coulomb slice is a true “infinitesimal slice” to the whole gauge group in $\mathcal{C}^\sigma(Y)$; i.e., a complement to the tangent space to the gauge orbits. Another such complement is the local Coulomb slice in the blow-up, \mathcal{K}_j^σ . The two slices are related as follows:

Lemma 5.5. *Let $x = (a, s, \phi) \in W_k^\sigma$. Then:*

(a) *The local Coulomb projection $\Pi_x^{\text{IC},\sigma}$ induces a linear isomorphism between the slices $\mathcal{K}_{j,x}^{\text{agC},\sigma}$ and $\mathcal{K}_{j,x}^\sigma$. Its inverse is the anticircular global Coulomb projection $\Pi^{\text{agC},\sigma}|_{\mathcal{K}_{j,x}^\sigma}$.*

(b) *If $s = 0$, then $\mathcal{K}_{j,x}^{\text{agC},\sigma} = \mathcal{K}_{j,x}^\sigma$, and $\Pi_x^{\text{agC},\sigma}|_{\mathcal{K}_{j,x}^\sigma} = (\Pi_*^{\text{gC},\sigma})_x|_{\mathcal{K}_{j,x}^\sigma} : \mathcal{K}_{j,x}^\sigma \rightarrow \mathcal{K}_{j,x}^{\text{agC},\sigma}$ is the identity.*

(c) *For any x , we have that $\Pi_x^{\text{agC},\sigma}|_{\mathcal{K}_{j,x}^\sigma} = (\Pi_*^{\text{gC},\sigma})_x|_{\mathcal{K}_{j,x}^\sigma}$.*

Proof. (a) As previously noted, both slices are complements to $\mathcal{J}_{j,x}^\sigma$. Further, both local Coulomb projection and anticircular global Coulomb projection are given by adding the suitable elements in $\mathcal{J}_{j,x}^\sigma$ that move to the other slice. This implies that the two maps are inverse to each other.

(b) If $s = 0$, we see from (66) and (68) that $\Pi_x^{\text{IC},\sigma} = (\Pi_*^{\text{gC},\sigma})_x$. The conclusion follows since $(\Pi_*^{\text{gC},\sigma})_x$ is both idempotent and invertible.

(c) The case $s = 0$ was studied in part (b). Therefore, it now suffices to consider the case when $s \neq 0$, so that x is irreducible. Then, the local slices are isomorphic to the corresponding ones in the blow-down, and the projections commute with the infinitesimal blow-down map, so we can simply work in the blow-down.

Let $\Phi = s\phi$ and $z = (a, \Phi)$. We need to check that if we have a vector $v \in \mathcal{K}_{j,z}$, then its projection $w = (\Pi_*^{\text{gC}})_z(v)$ lands in the anticircular global Coulomb slice. In other words, we know that $\text{Re}\langle(0, i\Phi), v\rangle_{L^2} = 0$, and we want to check that $\langle(0, i\Phi), w\rangle_{\tilde{g}} = 0$. Recall from (35) that Π_z^{elC} and $(\Pi_*^{\text{gC}})_z$ are inverse to each other; hence, $\Pi_z^{\text{elC}}(w) = v$. Using (37), we get

$$\langle(0, i\Phi), w\rangle_{\tilde{g}} = \text{Re}\langle(0, i\Phi), \Pi_z^{\text{elC}}(w)\rangle_{L^2} = \text{Re}\langle(0, i\Phi), v\rangle_{L^2} = 0,$$

as desired. \square

Recall that in Section 3.2 we defined an enlarged local Coulomb slice \mathcal{K}^e , complementary to the orbit of the normalized gauge group \mathcal{G}° . There is a similar enlarged local Coulomb slice in the blow-up,

$$\mathcal{K}_{j,x}^{e,\sigma} = \{(b, t, \psi) \in \mathcal{T}_{j,x}^\sigma \mid -d^*b + is^2 \text{Re}\langle i\phi, \psi \rangle^\circ = 0\}.$$

Note that the condition $-d^*b + is^2 \text{Re}\langle i\phi, \psi \rangle^\circ = 0$ simply means that $-d^*b + is^2 \text{Re}\langle i\phi, \psi \rangle$ is a constant function.

We define the enlarged local Coulomb projection on the blow-up, $\Pi_{(a,s,\phi)}^{\text{elC},\sigma} : \mathcal{T}_{j,x}^\sigma \rightarrow \mathcal{K}_{j,x}^{e,\sigma}$, by

$$(75) \quad \Pi_{(a,s,\phi)}^{\text{elC},\sigma}(b, r, \psi) = (b - d\zeta, r, \psi + \zeta\phi),$$

where ζ is such that $\int_Y \zeta = 0$ and $-d^*(b - d\zeta) + is^2 \text{Re}\langle i\phi, \psi + \zeta\phi \rangle^\circ = 0$.

Let

$$(76) \quad \mathcal{J}_{j,x}^{\circ,\sigma} = \{(-d\xi, 0, \xi\phi) \mid \int_Y \xi = 0\} \subset \mathcal{J}_{j,x}^\sigma$$

be the tangent to the orbit of \mathcal{G}° in the blow-up. We have direct sum decompositions

$$(77) \quad \mathcal{T}_{j,x}^\sigma = \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{T}_{j,x}^{\text{gC},\sigma}$$

and

$$(78) \quad \mathcal{T}_{j,x}^\sigma = \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{K}_{j,x}^{e,\sigma}.$$

Thus, $\mathcal{T}_{j,x}^{\text{gC},\sigma}$ and $\mathcal{K}_{j,x}^{e,\sigma}$ are both infinitesimal slices for the action of \mathcal{G}° .

Here is the analogue of Lemma 5.5; see also (35) for the corresponding result in the blow-down.

Lemma 5.6. *Let $x = (a, s, \phi) \in W_k^\sigma$. Then:*

- (a) *The enlarged local Coulomb projection $\Pi_x^{\text{elC},\sigma}$ induces a linear isomorphism between the slices $\mathcal{T}_{j,x}^{\text{gC},\sigma}$ and $\mathcal{K}_{j,x}^{\text{e},\sigma}$. Its inverse is the infinitesimal global Coulomb projection $\Pi_*^{\text{gC},\sigma}|_{\mathcal{K}_{j,x}^{\text{e},\sigma}}$.*
- (b) *If $s = 0$, then $\mathcal{T}_{j,x}^{\text{gC},\sigma} = \mathcal{K}_{j,x}^{\text{e},\sigma}$, and $(\Pi_*^{\text{gC},\sigma})_x|_{\mathcal{K}_{j,x}^{\text{e},\sigma}} : \mathcal{K}_{j,x}^{\text{e},\sigma} \rightarrow \mathcal{T}_{j,x}^{\text{gC},\sigma}$ is the identity.*

Finally, at $x \in W_k^\sigma$, let us introduce the *shear map*

$$S_x : \mathcal{T}_{j,x}^\sigma \rightarrow \mathcal{T}_{j,x}^\sigma$$

given by

$$(79) \quad \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{K}_{j,x}^{\text{e},\sigma} \rightarrow \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{T}_{j,x}^{\text{gC},\sigma}, \quad v \oplus w \mapsto v \oplus (\Pi_*^{\text{gC},\sigma})_x(w).$$

Its inverse is

$$S_x^{-1} : \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{T}_{j,x}^{\text{gC},\sigma} \rightarrow \mathcal{J}_{j,x}^{\circ,\sigma} \oplus \mathcal{K}_{j,x}^{\text{e},\sigma}, \quad v \oplus w \mapsto v \oplus \Pi_x^{\text{elC},\sigma}(w).$$

We can write, in a more compressed form,

$$\begin{aligned} S_x(v) &= v - \Pi_x^{\text{elC},\sigma}(v) + (\Pi_*^{\text{gC},\sigma})_x(v), \\ S_x^{-1}(v) &= v + \Pi_x^{\text{elC},\sigma}(v) - (\Pi_*^{\text{gC},\sigma})_x(v). \end{aligned}$$

Putting together all the maps S_x , we obtain an automorphism S of the bundle \mathcal{T}_j^σ over W_k^σ .

Remark 5.7. With regard to the infinitesimal slices to the whole gauge action, observe that S_x maps $\mathcal{K}_{j,x}^\sigma$ to $\mathcal{K}_{j,x}^{\text{agC},\sigma}$; compare Lemma 5.5. Also, S_x preserves (but does not act by the identity on) the infinitesimal orbit space $\mathcal{J}_{j,x}^\sigma$.

5.2. Choices of gauge on cylinders. Let $Z = I \times Y$ be a cylinder. Recall that $\mathcal{C}(Z)$ consists of pairs (a, ϕ) with $a \in \Omega^1(Z; i\mathbb{R})$ and $\phi \in \Gamma(\mathbb{S}^+)$. We write such a pair as a path

$$(a(t) + \alpha(t)dt, \phi(t)), \quad t \in I,$$

where $a(t) \in \Omega^1(Y; i\mathbb{R})$, $\alpha(t) \in C^\infty(Y; i\mathbb{R})$, and $\phi(t) \in \Gamma(\mathbb{S})$.

Recall that if $\alpha(t) = 0$ we say that (a, ϕ) is in temporal gauge and that any configuration can be put into temporal gauge using the action of $\mathcal{G}(Z)$.

We seek an analogue of $\mathcal{C}(Z)$ adapted to global Coulomb gauge. The first guess is to consider $W(Z)$, the subspace of $\mathcal{C}(Z)$ consisting of configurations in temporal gauge, and that are also in (three-dimensional) Coulomb gauge at each slice $\{t\} \times Y$; that is,

$$W(Z) = \{(a, \phi) \in \Gamma(Z, p^*(iT^*Y \oplus \mathbb{S})) \mid a(t) \in \ker d^*, \forall t\},$$

where $p : Z \rightarrow Y$ denotes the projection. Note that an arbitrary configuration $(a, \phi) \in \mathcal{C}(Z)$ cannot always be moved into $W(Z)$ by a four-dimensional gauge transformation; instead, we can move it into temporal gauge, and then Coulomb project in each three-dimensional slice. Further, on $W(Z)$ we could consider an action by the group of slicewise constant gauge transformations

$$\mathcal{G}^{\text{gC}}(Z) = C^\infty(I; S^1),$$

with $u \in \mathcal{G}^{\text{gC}}(Z)$ acting by $(a(t), \phi(t)) \rightarrow (a(t), u(t)\phi(t))$.

As we saw in Section 3.3, the global Coulomb projections of Seiberg-Witten trajectories are the solutions of

$$(80) \quad \left(\frac{d}{dt} + \mathcal{X}^{\text{gC}} \right) (a(t), \phi(t)) = 0,$$

where $(a, \phi) \in W(Z)$.

Note that the equations (80) are invariant under the action of constant $u \in S^1$, but not under all of $\mathcal{G}^{\text{gC}}(Z)$. This is similar to what happens in $\mathcal{C}(Z)$: the Seiberg-Witten equations are invariant under the whole group $\mathcal{G}(Z)$, but once we move to temporal gauge and write the equations as a gradient flow, we are only left with the action of $\mathcal{G}(Y)$, constant in t .

In view of this, a better analogue of $\mathcal{C}(Z)$ is defined using the following notion:

Definition 5.8. *A configuration $(a(t) + \alpha(t)dt, \phi(t)) \in \Omega^1(Z; i\mathbb{R}) \oplus \Gamma(\mathbb{S}^+)$ is said to be in pseudo-temporal gauge if for each t , the component $\alpha(t)$ is constant as a function on Y .*

We let $\mathcal{C}^{\text{gC}}(Z)$ consist of pairs $(a(t) + \alpha(t)dt, \phi(t))$ in pseudo-temporal gauge, and such that $(a(t), \phi(t))$ is in slice-wise Coulomb gauge, i.e.:

$$d(\alpha(t)) = 0, \quad d^*(a(t)) = 0, \quad \forall t \in I.$$

The elements of $\mathcal{G}^{\text{gC}}(Z)$ act on $\mathcal{C}^{\text{gC}}(Z)$ as usual gauge transformations:

$$u : (a(t) + \alpha(t)dt, \phi(t)) \mapsto (a(t) + (\alpha(t) - u(t)^{-1} \frac{du(t)}{dt})dt, u(t)\phi(t)).$$

Consider the process of moving an arbitrary configuration in $\mathcal{C}(Z)$ into pseudo-temporal gauge by an element of $\mathcal{G}(Z)$, and then applying slice-wise global Coulomb projection to land in $\mathcal{C}^{\text{gC}}(Z)$. Under this process, the Seiberg-Witten equations turn into:

$$(81) \quad \left(\frac{d}{dt} + \mathcal{X}^{\text{gC}} \right) (a(t), \phi(t)) + (0, \alpha(t)\phi(t)) = 0,$$

for $(a(t) + \alpha(t)dt, \phi(t)) \in \mathcal{C}^{\text{gC}}(Z)$. These equations are invariant under the action of $\mathcal{G}^{\text{gC}}(Z)$.

Of course, every solution of (81) can be transformed into a solution of (80) by moving into temporal gauge. Most of the time we will work with solutions of (80), i.e., trajectories of \mathcal{X}^{gC} . However, when considering infinitesimal deformations of such trajectories (as we shall do in Section 5.9, for example), it is important to allow the more general pseudo-temporal gauge in order to obtain a good Fredholm problem.

By analogy with the section $\mathcal{F} : \mathcal{C}(Z) \rightarrow \mathcal{V}(Z)$ from Section 4.1, we write $\mathcal{F}^{\text{gC}}(a, \phi)$ for the left hand side of Equation (81). We view \mathcal{F}^{gC} as a section

$$\mathcal{F}^{\text{gC}} : \mathcal{C}^{\text{gC}}(Z) \rightarrow \mathcal{V}^{\text{gC}}(Z),$$

where $\mathcal{V}^{\text{gC}}(Z)$ is the trivial $W(Z)$ bundle over $\mathcal{C}^{\text{gC}}(Z)$.

It is worth comparing $\mathcal{V}^{\text{gC}}(Z)$ to the bundle $\mathcal{V}(Z)$ from Section 4.1, whose fibers were $\Gamma(Z; i\Lambda_+^2 T^*Z \oplus \mathbb{S}^-)$. In our setting, we identify self-dual imaginary two-forms on Z with sections $a = (a(t))$ of $p^*(iT^*Y)$, via sending a section a to $*a + dt \wedge a$, where $*$ here denotes the Hodge star operator on Y . We then impose a Coulomb gauge condition $d^*a(t) = 0$ for all t . Further, the bundle \mathbb{S}^- can be identified with $p^*\mathbb{S}$. Thus, the fibers become $W(Z)$, as in our definition of $\mathcal{V}^{\text{gC}}(Z)$.

If Z is a compact cylinder, then starting from the space $W(Z)$ we can consider Sobolev completions $W_k(Z)$ and blow-ups $W_k^\sigma(Z)$ and $W_k^\tau(Z)$. The space $W_k^\tau(Z)$ is a subset of the Banach manifold $\widetilde{W}_k^\tau(Z)$, the latter of which is obtained by removing the condition $s(t) \geq 0$;

compare with (46). Similarly, we can define $\mathcal{C}_k^{\text{gC}}(Z)$, its σ and τ blow-ups, and a Banach manifold $\widetilde{\mathcal{C}}_k^{\text{gC},\tau}(Z)$.

The tangent space to $W_k(Z)$ can be completed to $T_j W(Z)$ (for $j \leq k$), which is a trivial bundle with fiber $W_j(Z)$. The tangent bundle to $\mathcal{C}_k^{\text{gC}}(Z)$ is denoted $\mathcal{T}_k^{\text{gC}}(Z)$, and has completions $\mathcal{T}_j^{\text{gC}}(Z)$. There are blown-up analogues $\mathcal{T}_j^{\text{gC},\sigma}(Z)$ and $\mathcal{T}_j^{\text{gC},\tau}(Z)$, the latter of which we think of as a bundle over $\widetilde{\mathcal{C}}_k^{\text{gC},\tau}(Z)$. Similar constructions can be done for the bundle $\mathcal{V}(Z)$.

If Z is an infinite cylinder, we will also consider $L_{j,\text{loc}}^2$ completions. These are denoted $W_{k,\text{loc}}(Z)$, $\mathcal{C}_{k,\text{loc}}^{\text{gC}}(Z)$, and so on.

5.3. Controlled Coulomb perturbations. In Section 4.2 we discussed how one can perturb the Seiberg-Witten equations by a formal gradient \mathfrak{q} . In Coulomb gauge, a perturbation is an S^1 -equivariant vector field

$$\eta : W \rightarrow \mathcal{T}_0^{\text{gC}}.$$

We write η^0 and η^1 for the connection and spinor components of η .

We are interested in perturbations η that are \tilde{g} -formal gradients of functions $f : W \rightarrow \mathbb{R}$. Given such an η , by applying it slicewise in t on a cylinder $Z = I \times Y$, we obtain a section

$$(82) \quad \hat{\eta} : \mathcal{C}^{\text{gC}}(Z) \rightarrow \mathcal{V}^{\text{gC}}(Z), \quad \hat{\eta}(a(t) + \alpha(t)dt, \phi(t)) = \eta(a(t), \phi(t)).$$

We require that $\hat{\eta}$ preserves Sobolev regularity. We will also need properties of η analogous to some of the very tameness properties of \mathfrak{q} . (Compare Definitions 4.5 and 4.9.)

Definition 5.9. *Suppose that $\eta = (\eta^0, \eta^1)$ is the \tilde{g} -formal gradient of an S^1 -equivariant function $f : W \rightarrow \mathbb{R}$. Let $\hat{\eta}$ be the induced four-dimensional perturbation. We say η is a controlled Coulomb perturbation if for all integers $k \geq 2$, and all compact cylinders $Z = [t_1, t_2] \times Y$,*

- (i) $\hat{\eta}$ defines an element of $C_{\text{fb}}^\infty(\mathcal{C}_k^{\text{gC}}(Z), \mathcal{V}_k^{\text{gC}}(Z))$;
- (ii) The first derivative $\mathcal{D}\hat{\eta} \in C_{\text{fb}}^\infty(\mathcal{C}_k^{\text{gC}}(Z), \text{Hom}(\mathcal{T}_k^{\text{gC}}(Z), \mathcal{V}_k^{\text{gC}}(Z)))$ extends to a C_{fb}^∞ map into $\text{Hom}(\mathcal{T}_j^{\text{gC}}(Z), \mathcal{V}_j^{\text{gC}}(Z))$, for all integers $j \in [-k, k]$;

Remark 5.10. In [KM07], Properties (ii), (iv), (v), and (vi) in Definition 4.5 are used to get bounds on the stationary points and trajectories of $\mathcal{X}_{\mathfrak{q}}$. When we do finite-dimensional approximation, we have a priori bounds since we restrict to the ball $B(2R)$. Therefore, there is no need to include analogues of these properties in Definition 5.9 above.

We recall that the infinitesimal global Coulomb projection of the Seiberg-Witten vector field $\mathcal{X} = \text{grad } \mathcal{L}$ is given by $\mathcal{X}^{\text{gC}} = l + c$, where l and c are defined by (39) and (40) respectively. Suppose that \mathfrak{q} is an abstract perturbation given by the formal gradient of a function f . Let

$$(83) \quad \eta_{\mathfrak{q}}(a, \phi) = (\Pi_*^{\text{gC}})_{(a, \phi)} \mathfrak{q}(a, \phi)$$

be the infinitesimal global Coulomb projection of \mathfrak{q} . The infinitesimal global Coulomb projection of $\mathcal{X}_{\mathfrak{q}} = \mathcal{X} + \mathfrak{q}$ is then

$$\mathcal{X}_{\mathfrak{q}}^{\text{gC}} = \mathcal{X}^{\text{gC}} + \eta_{\mathfrak{q}},$$

which is the gradient of $(\mathcal{L} + f)|_W$ with respect to the metric \tilde{g} introduced in (36).

Lemma 5.11. *Let \mathfrak{q} be a very tame perturbation. Then, $\eta_{\mathfrak{q}}$ is a controlled Coulomb perturbation.*

Before proving the claim, we will need to prove some additional properties of the infinitesimal global Coulomb projection.

Lemma 5.12. *Fix $j \in [-k, k]$. Then,*

- (a) *the map $\Pi_*^{\text{gC}} : \mathcal{T}_j \rightarrow \mathcal{T}_j^{\text{gC}}$ is functionally bounded as a map from $\mathcal{C}_k(Y)$ to $\text{Hom}(\mathcal{C}_j(Y), W_j)$;*
- (b) *if $X \in C_{\text{fb}}^n(W_k, \mathcal{C}_k(Y))$ satisfies that $\mathcal{D}X : W_k \rightarrow \text{Hom}(W_k, \mathcal{C}_k(Y))$ extends to an element of $C_{\text{fb}}^{n-1}(W_k, \text{Hom}(W_j, \mathcal{C}_j(Y)))$, then $\mathcal{D}(\Pi_*^{\text{gC}} \circ X) : W_k \rightarrow \text{Hom}(W_j, W_j)$ is also in C_{fb}^{n-1} ;*
- (c) *the analogues of (a) and (b) hold for $I \times Y$ as well for $I \subseteq \mathbb{R}$ a closed interval (with Π_*^{gC} applied slice-wise).*

Proof. (a) Let $(a, \phi) \in \mathcal{C}_k(Y)$. By definition,

$$(\Pi_*^{\text{gC}})_{(a, \phi)}(b, \psi) = (\pi(b), e^{Gd^*a}((Gd^*b)\phi + \psi)).$$

We are interested in showing that this expression has L_j^2 norm bounded in terms of the L_k^2 norm of (a, ϕ) and the L_j^2 norm of (b, ψ) . This is clear for the first term $\pi(b)$, because π is a linear, continuous operator taking L_j^2 forms to L_j^2 forms. For the second term, the continuity and bilinearity of the Sobolev multiplication $L_j^2 \times L_k^2 \rightarrow L_j^2$ gives bounds

$$\|(Gd^*b)\phi\|_{L_j^2} \leq C\|Gd^*b\|_{L_j^2}\|\phi\|_{L_k^2}.$$

We then use the linearity and continuity of Gd^* to bound $\|Gd^*b\|_{L_j^2}$ in terms of $\|b\|_{L_j^2}$. The functional boundedness of the exponential map gives a bound on e^{Gd^*a} . Using the Sobolev multiplication again, we get the desired L_j^2 bounds on $e^{Gd^*a}((Gd^*b)\phi + \psi)$.

(b) By Lemma 3.1, we have that Π_*^{gC} is smooth and therefore we are only interested in functional boundedness. Precisely, to show that $\mathcal{D}(\Pi_*^{\text{gC}} \circ X)$ is in C_{fb}^{m-1} , we need to check that $\mathcal{D}^m(\Pi_*^{\text{gC}} \circ X)$ is functionally bounded for all $m = 1, \dots, n$.

We first consider the case $m = 1$. Let $(a, \phi) \in W_k$ and $(b, \psi) \in W_j$. It is straightforward to compute

$$(84) \quad \mathcal{D}_{(a, \phi)}(\Pi_*^{\text{gC}} \circ X)(b, \psi) = \left(\pi(\mathcal{D}_{(a, \phi)}X^0(b, \psi)), \mathcal{D}_{(a, \phi)}X^1(b, \psi) + (Gd^*\mathcal{D}_{(a, \phi)}X^0(b, \psi))\phi + (Gd^*X^0(a, \phi))\psi \right).$$

Since X and $\mathcal{D}X$ are functionally bounded (the latter as a map from W_k to $\text{Hom}(W_j, \mathcal{C}_j(Y))$), we get bounds on $X(a, \phi)$ and $\mathcal{D}_{(a, \phi)}X^0(b, \psi)$ in terms of the L_k^2 norm of (a, ϕ) and the L_j^2 norm of (b, ψ) . To obtain the desired L_j^2 bounds on (84), as above, we apply the continuity and bilinearity of the Sobolev multiplication $L_k^2 \times L_j^2 \rightarrow L_j^2$.

Finally, we give the argument for the second derivative (as this will illustrate the appropriate Sobolev norms) and allow the reader to complete the proof by induction. We consider the derivative of $\mathcal{D}(\Pi_*^{\text{gC}} \circ X)$, which we think of as a map from W_k to $\text{Hom}(W_k \times W_j, W_j)$. Again, let $(a, \phi) \in W_k$, $(b, \psi) \in W_j$. We also denote an L_k^2 tangent vector to (a, ϕ) by (α, ζ) .

Direct computation shows

$$\begin{aligned} \mathcal{D}_{(a,\phi)}^2(\Pi_*^{\text{gC}} \circ X)((\alpha, \zeta), (b, \psi)) &= \left(\pi(\mathcal{D}_{(a,\phi)}^2 X^0)(\alpha, \zeta), (b, \psi) \right), \\ &\quad (Gd^*(\mathcal{D}_{(a,\phi)}^2 X^0)((\alpha, \zeta), (b, \psi)))\phi + Gd^*(\mathcal{D}_{(a,\phi)} X^0(b, \psi))\zeta \\ &\quad + (\mathcal{D}_{(a,\phi)}^2 X^1)(\alpha, \zeta), (b, \psi) \Big). \end{aligned}$$

Again, the functional boundedness of X and $\mathcal{D}X$ together with Sobolev multiplication give the desired result.

(c) Similar arguments apply to establish the four-dimensional analogues. \square

Proof of Lemma 5.11. First, note that if \mathfrak{q} is the L^2 -formal gradient of $f : \mathcal{C}(Y) \rightarrow \mathbb{R}$, then $\eta_{\mathfrak{q}}$ is the \tilde{g} -formal gradient of the restriction of f to W .

Since $\eta_{\mathfrak{q}}(a, \phi) = (\Pi_*^{\text{gC}})_{(a,\phi)} \mathfrak{q}(a, \phi)$, the result follows by combining Lemmas 3.1 and 5.12 with the corresponding Properties (i) and (iii) of \mathfrak{q} in Definition 4.9, since \mathfrak{q} is a very tame perturbation. \square

5.4. Trajectories in global Coulomb gauge. Let \mathfrak{q} be a very tame perturbation and $\eta_{\mathfrak{q}}$ the induced controlled perturbation in Coulomb gauge.

Let us write $\mathcal{X}_{\mathfrak{q}}^{\text{gC}} = ((\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^0, (\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^1)$ where, as usual, the superscript 0 denotes the form part and the superscript 1 denotes the spinorial part. The vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$ on W induces a vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ on W^{σ} , given by

$$(85) \quad \mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}(a, s, \phi) = ((\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^0(a, s\phi), \Lambda_{\mathfrak{q}}(a, s, \phi)s, (\widetilde{(\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^1}(a, s, \phi) - \Lambda_{\mathfrak{q}}(a, s, \phi)\phi),$$

where

$$(86) \quad (\widetilde{(\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^1}(a, s, \phi) = \int_0^1 \mathcal{D}_{(a, sr\phi)}(\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^1(0, \phi) dr,$$

and

$$(87) \quad \Lambda_{\mathfrak{q}}(a, s, \phi) = \text{Re}\langle \phi, (\widetilde{(\mathcal{X}_{\mathfrak{q}}^{\text{gC}})^1}(a, s, \phi) \rangle_{L^2}.$$

Since $\mathcal{X}_{\mathfrak{q}}^{\text{gC}} = \Pi_*^{\text{gC}} \circ \mathcal{X}_{\mathfrak{q}}$, we have that $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma} = \Pi_*^{\text{gC},\sigma} \circ \mathcal{X}_{\mathfrak{q}}^{\sigma}$.

Lemma 5.13. *The vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma} : W_k^{\sigma} \rightarrow \mathcal{T}_{k-1}^{\text{gC},\sigma}$ is smooth.*

Proof. We write $\mathcal{X}_{\mathfrak{q}}^{\text{gC}} = \Pi_*^{\text{gC}} \circ (\mathcal{X} + \mathfrak{q})$. Lemma 3.1 shows that $\Pi_*^{\text{gC}} : \mathcal{T}_{k-1} \rightarrow \mathcal{T}_{k-1}^{\text{gC}}$ is smooth. Since $\mathcal{X} = \text{grad } \mathcal{L}$ and \mathfrak{q} are smooth as maps from $\mathcal{C}_k(Y)$ to \mathcal{T}_{k-1} , we get that $\mathcal{X}_{\mathfrak{q}}^{\text{gC}} : W_k \rightarrow \mathcal{T}_{k-1}^{\text{gC}}$ is smooth. Hence, the induced vector field on the blow-up, $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$, is also smooth (compare [KM07, Lemma 10.2.1]). \square

We are interested in the dynamics of the vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ on W^{σ} .

Every stationary point of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ on $\mathcal{C}^{\sigma}(Y)$ can be moved into W^{σ} by the global Coulomb projection. Conversely, every stationary point x of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ on W^{σ} is also a stationary point of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$, since $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ lands in \mathcal{K}_x^{σ} and infinitesimal global Coulomb projection induces an isomorphism from \mathcal{K}_x^{σ} to $\mathcal{K}_x^{\text{agC},\sigma}$. Thus, $\Pi^{\text{gC},\sigma}$ induces a bijection

$$(88) \quad \{\text{stationary points of } \mathcal{X}_{\mathfrak{q}}^{\sigma}\} / \mathcal{G} \xrightarrow{\cong} \{\text{stationary points of } \mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}\} / S^1.$$

Note that the condition of being irreducible, boundary stable, or boundary unstable is preserved by this bijection. By contrast, a trajectory⁹ of $\mathcal{X}_q^{\text{gC},\sigma}$ on W^σ is *not* a trajectory of \mathcal{X}_q^σ ; still, we have:

Proposition 5.14. *The trajectories of $\mathcal{X}_q^{\text{gC},\sigma}$ on W^σ are precisely the global Coulomb projections of the trajectories of \mathcal{X}_q^σ on the blow-up $\mathcal{C}^\sigma(Y)$. In fact, global Coulomb projection induces a bijection*

$$(89) \quad \{\text{trajectories of } \mathcal{X}_q^\sigma\} / \mathcal{G} \xrightarrow{\cong} \{\text{trajectories of } \mathcal{X}_q^{\text{gC},\sigma}\} / S^1,$$

where \mathcal{G} acts on trajectories $x(t)$ by three-dimensional gauge transformations, constant in t .

Remark 5.15. Because q is a very tame perturbation, every L_k^2 trajectory of \mathcal{X}_q^σ is gauge-equivalent to a smooth one. Furthermore, every trajectory of $\mathcal{X}_q^{\text{gC},\sigma}$ is smooth.

Proof of Proposition 5.14. If γ is a smooth trajectory of \mathcal{X}_q^σ , its global Coulomb projection γ^\flat is a trajectory of $\Pi_*^{\text{gC},\sigma} \circ \mathcal{X}_q^\sigma$, which is exactly $\mathcal{X}_q^{\text{gC},\sigma}$.

Conversely, given a trajectory γ of $\mathcal{X}_q^{\text{gC},\sigma}$, we can lift it to a smooth trajectory γ^\sharp of \mathcal{X}_q^σ as follows. Consider the submanifold $\mathcal{O}(\gamma)_k \subset \mathcal{C}_k^\sigma(Y)$ consisting of all points that are on the gauge orbits of points on γ . Thus, any $x \in \mathcal{O}(\gamma)_k$ can be written as $x = u \cdot \gamma(t_0)$ where u is an L_{k+1}^2 gauge transformation on Y and $t_0 \in \mathbb{R}$. The vector $v_x = \mathcal{X}_q^\sigma(x)$ is in the local Coulomb slice $\mathcal{K}_{x,k-1}^\sigma$. Consider the push-forward $(u^{-1})_* v_x$, which is in the local Coulomb slice to $u^{-1} \cdot x = \Pi_*^{\text{gC}}(x)$. Now, $u^{-1} \cdot x$ is part of the trajectory γ , and hence is smooth. Therefore, $(u^{-1})_* v_x = (\mathcal{X}_q^\sigma)_{u^{-1} \cdot x}$ is also smooth. Pushing it back by the L_{k+1}^2 gauge transformation u yields the vector v , which we now see that it must be in L_k^2 . Thus, the vectors v_x form a true vector field on the Banach manifold $\mathcal{O}(\gamma)_k$. By integrating this vector field starting at a smooth point $x \in \mathcal{O}(\gamma)_k$, we obtain a lift γ^\sharp of γ . We can do this for any k , and obtain the same lift; hence, γ^\sharp is smooth.

In the above construction, note that the lift γ^\sharp is unique up to transformation by an element in \mathcal{G} . Indeed, in four-dimensions, after moving to temporal gauge, there are only gauge transformations which are constant in t . (Compare [KM07, Proposition 7.2.1].) \square

Using the formula (67) for $(\Pi_*^{\text{gC},\sigma})_{(a,s,\phi)}$, we can describe the trajectories of $\mathcal{X}_q^{\text{gC},\sigma}$ more explicitly, as the paths $(a(t), s(t), \phi(t))$ in slicewise Coulomb gauge that satisfy

$$(90) \quad \begin{aligned} \frac{d}{dt} a &= - * da - s^2 \pi(\tau(\phi, \phi)) - \pi(q^0(a, s\phi)), \\ \frac{d}{dt} s &= -\Lambda_q(a, s, \phi) s, \\ \frac{d}{dt} \phi &= -D_a \phi - \tilde{q}^1(a, s, \phi) + \Lambda_q(a, s, \phi) \phi - s^2 Gd^*(\tau(\phi, \phi)) \phi - Gd^*(q^0(a, s\phi)) \phi. \end{aligned}$$

⁹From now on, by trajectory we will always mean a trajectory of finite type as in Section 3.3.

Alternatively, we can write the trajectory equations as

$$\begin{aligned}
 (91) \quad & \frac{d}{dt}a = - * da - s^2 \pi(\tau(\phi, \phi)) - \pi(\mathbf{q}^0(a, s\phi)), \\
 & \frac{d}{dt}s = -\Lambda_{\mathbf{q}}(a, s, \phi)s, \\
 & \frac{d}{dt}\phi = -D_a\phi - s^2\xi(\phi)\phi - \tilde{\eta}_{\mathbf{q}}^1(a, s, \phi) + \Lambda_{\mathbf{q}}(a, s, \phi)\phi \\
 & \quad = -D\phi - \tilde{c}^1(a, s, \phi) - \tilde{\eta}_{\mathbf{q}}^1(a, s, \phi) + \Lambda_{\mathbf{q}}(a, s, \phi)\phi,
 \end{aligned}$$

where $\xi(\phi)$ is as in (40), and we use $\tilde{X}^1(a, s, \phi)$ to denote $\int_0^1 \mathcal{D}_{(a, sr\phi)} X^1(0, \phi) dr$.

By making use of $\mathcal{G}^{\text{SC}}(Z)$ -equivariance as in (82), the equations defining the flow of $\mathcal{X}_{\mathbf{q}}^{\text{SC}, \sigma}$ define a section

$$(92) \quad \mathcal{F}_{\mathbf{q}}^{\text{SC}, \tau} : \mathcal{C}^{\text{SC}, \tau}(Z) \rightarrow \mathcal{V}^{\text{SC}, \tau}(Z),$$

so that in temporal gauge we have

$$\mathcal{F}_{\mathbf{q}}^{\text{SC}, \tau} = \frac{d}{dt} + \mathcal{X}_{\mathbf{q}}^{\text{SC}, \sigma}.$$

It follows that the space of trajectories of $\mathcal{X}_{\mathbf{q}}^{\text{SC}, \sigma}$ modulo the action of S^1 is the zero set of $\mathcal{F}_{\mathbf{q}}^{\text{SC}, \tau}$ modulo the action of $\mathcal{G}^{\text{SC}}(Z)$. As in Section 4.5, it will be useful to consider the extension of $\mathcal{F}_{\mathbf{q}}^{\text{SC}, \tau}$ to $\tilde{\mathcal{C}}^{\text{SC}, \tau}(Z)$, so that we may differentiate this map.

5.5. Hessians. In Section 12.4 in [KM07], Kronheimer and Mrowka study the derivative of the vector field $\mathcal{X}_{\mathbf{q}}^{\sigma}$ on $\mathcal{C}_k^{\sigma}(Y)$. At each $x \in \mathcal{C}_k^{\sigma}(Y)$, this is an operator between the local Coulomb slices:

$$\text{Hess}_{\mathbf{q}, x}^{\sigma} : \mathcal{K}_{k, x}^{\sigma} \rightarrow \mathcal{K}_{k-1, x}^{\sigma},$$

called the *Hessian* (in the blow-up). The Hessian is needed in Floer theory for several reasons. First, the non-degeneracy of a stationary point x is expressed in terms of the surjectivity of $\text{Hess}_{\mathbf{q}, x}^{\sigma}$. Second, we need to study Hessians at all points (not necessarily stationary) in order to construct the operator

$$(93) \quad \frac{d}{dt} + \text{Hess}_{\mathbf{q}}^{\sigma},$$

which can be applied to vector fields along a path γ in $\mathcal{C}^{\sigma}(Y)$. If γ is a flow trajectory for $\mathcal{X}_{\mathbf{q}}^{\sigma}$, then the surjectivity of $d/dt + \text{Hess}_{\mathbf{q}}^{\sigma}$ indicates that the moduli space of trajectories is regular at γ . Also, for arbitrary γ (not necessarily a flow trajectory), we need the operator (93) to describe the relative grading of stationary points, and to construct orientations on the moduli spaces of trajectories.

We will do the same analysis in global Coulomb gauge. We will first define a \tilde{g} -Hessian before the blow-up, at irreducible points (a, ϕ) (that is, those with $\phi \neq 0$). We will then construct a \tilde{g} -Hessian on the blow-up, well-defined at all points.

5.5.1. The Hessian at irreducibles. We start by recalling the original Hessian at irreducibles, as defined in Section 12.3 of [KM07]:

$$(94) \quad \text{Hess}_{\mathbf{q}, x} = \Pi_x^{\text{IC}} \circ \mathcal{D}_x \mathcal{X}_{\mathbf{q}} : \mathcal{K}_{k, x} \rightarrow \mathcal{K}_{k-1, x}.$$

This formula applies to any $x = (a, \phi) \in \mathcal{C}_k(Y)$ with $\phi \neq 0$.

Remark 5.16. When x is a stationary point, we can actually drop the projection Π_x^{IC} from the formula for $\text{Hess}_{\mathbf{q},x}$. This is because $\mathcal{X}_{\mathbf{q}}$ is a formal gradient, so it is orthogonal to the gauge orbits; therefore, it can be viewed as a section of \mathcal{K}_{k-1} , and its derivative at a zero point must automatically land in the local Coulomb slice.

The \tilde{g} -Hessian is defined as:

$$(95) \quad \text{Hess}_{\mathbf{q},x}^{\tilde{g}} = \Pi_x^{\text{agC}} \circ \mathcal{D}_x^{\tilde{g}} \mathcal{X}_{\mathbf{q}}^{\text{gC}} : \mathcal{K}_{k,x}^{\text{agC}} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC}}.$$

Here, $\mathcal{D}^{\tilde{g}}$ denotes the connection on \mathcal{T}^{gC} induced by the \tilde{g} metric on W_k . To put this in context, recall that the derivative of a vector field is well-defined at a stationary point x , but when x is not stationary it depends on the choice of a connection in the tangent bundle. In (94), we simply viewed the vector field $\mathcal{X}_{\mathbf{q}}$ as a map into affine space, and took its derivative as such. Doing the same thing in global Coulomb gauge would mean using the connection \mathcal{D} induced by the ordinary L^2 metric. However, $\mathcal{X}_{\mathbf{q}}^{\text{gC}}$ is the \tilde{g} -gradient of $(\mathcal{L} + f)|_W$, and taking its ordinary derivative $\mathcal{D}_x \mathcal{X}_{\mathbf{q}}^{\text{gC}}$ would yield an operator that is not symmetric. It is more natural to use the connection $\mathcal{D}^{\tilde{g}}$, which is explicitly given by the formula:

$$(96) \quad \mathcal{D}^{\tilde{g}}(X) = \Pi_*^{\text{C}} \circ \mathcal{D}(\Pi^{\text{elC}}(X)) \circ \Pi^{\text{elC}}.$$

Here, X is a vector field on W_k , and $\Pi^{\text{elC}}(X)$ is a section of \mathcal{T}_k over W_k . We implicitly extend $\Pi^{\text{elC}}(X)$ to a \mathcal{G}_{k+1}° -invariant vector field on the whole configuration space $\mathcal{C}_k(Y)$; this allows us to take its derivative in the direction $\Pi^{\text{elC}}(Z)$, for some Z .

Let us check that the formula (96) produces a connection compatible with \tilde{g} . If X, Y and Z are vector fields on W_k , we have:

$$\begin{aligned} Z\langle X, Y \rangle_{\tilde{g}} &= Z \text{Re}\langle \Pi^{\text{elC}}(X), \Pi^{\text{elC}}(Y) \rangle_{L^2} \\ &= \Pi^{\text{elC}}(Z) \text{Re}\langle \Pi^{\text{elC}}(X), \Pi^{\text{elC}}(Y) \rangle_{L^2} \\ &= \text{Re}\langle (\mathcal{D}(\Pi^{\text{elC}}(X)))(\Pi^{\text{elC}}(Z)), \Pi^{\text{elC}}(Y) \rangle_{L^2} + \langle \Pi^{\text{elC}}(X), (\mathcal{D}(\Pi^{\text{elC}}(Y)))(\Pi^{\text{elC}}(Z)) \rangle_{L^2} \\ &= \text{Re}\langle (\mathcal{D}^{\tilde{g}}X)(Z), \Pi^{\text{elC}}(Y) \rangle_{L^2} + \langle \Pi^{\text{elC}}(X), (\mathcal{D}^{\tilde{g}}Y)(Z) \rangle_{L^2} \\ &= \text{Re}\langle \Pi^{\text{elC}}((\mathcal{D}^{\tilde{g}}X)(Z)), \Pi^{\text{elC}}(Y) \rangle_{L^2} + \langle \Pi^{\text{elC}}(X), \Pi^{\text{elC}}((\mathcal{D}^{\tilde{g}}Y)(Z)) \rangle_{L^2} \\ &= \text{Re}\langle (\mathcal{D}^{\tilde{g}}X)(Z), Y \rangle_{\tilde{g}} + \langle X, (\mathcal{D}^{\tilde{g}}Y)(Z) \rangle_{\tilde{g}}. \end{aligned}$$

Here, we have implicitly extended $\Pi^{\text{elC}}(X)$ and $\Pi^{\text{elC}}(Y)$ to \mathcal{G}_{k+1}° -invariant vector fields on the configuration space. The second equality above is true because the function

$$\text{Re}\langle \Pi^{\text{elC}}(X), \Pi^{\text{elC}}(Y) \rangle_{L^2}$$

is \mathcal{G}_{k+1}° -invariant, and therefore its derivative in the direction $Z - \Pi^{\text{elC}}(Z) \in \mathcal{J}_k^\circ$ vanishes. For some of the other equalities, we used repeatedly the fact that the L^2 inner product with something in the enlarged local Coulomb slice is unchanged by applying either Π_*^{gC} or Π^{elC} to the other factor. This last fact is true because both of these operations consist in adding a vector in \mathcal{J}_k° , and such vectors are L^2 perpendicular to the enlarged local Coulomb slice.

One can also check that the connection $\mathcal{D}^{\tilde{g}}$ is torsion-free:

$$\begin{aligned} (\mathcal{D}^{\tilde{g}}X)(Y) - (\mathcal{D}^{\tilde{g}}Y)(X) &= \Pi_*^{\text{gC}} \circ \mathcal{D}(\Pi^{\text{elC}}(X))(\Pi^{\text{elC}}(Y)) - \Pi_*^{\text{gC}} \circ \mathcal{D}(\Pi^{\text{elC}}(Y))(\Pi^{\text{elC}}(X)) \\ &= \Pi_*^{\text{gC}}[\Pi^{\text{elC}}(X), \Pi^{\text{elC}}(Y)] \\ &= [X, Y]. \end{aligned}$$

To see the last equality, it suffices to check how the respective vector fields act on \mathcal{G}_{k+1}° -invariant functions on the configuration space. On such functions f , the action of a vector field does not change if we apply Π^{elC} or Π_*^{gC} to that vector field (because this means changing it by a vector field in a direction tangent to \mathcal{G}_{k+1}°). The equality follows easily from this observation.

Remark 5.17. It follows from the above discussion that $\mathcal{D}^{\tilde{g}}$ is the exact analogue of the Levi-Civita connection on the infinite-dimensional manifold W with respect to the \tilde{g} -metric. Therefore, when we compute the \tilde{g} -Hessian of a real-valued function on W , we will obtain a symmetric operator.

From Section 3.2 we know that the projection Π_*^{gC} maps \mathcal{K}^e isomorphically onto \mathcal{T}^{gC} , with inverse Π^{elC} . Since $\mathcal{X}_q^{\text{gC}} = \Pi_*^{\text{gC}} \mathcal{X}_q$, we also have $\mathcal{X}_q = \Pi^{\text{elC}} \mathcal{X}_q^{\text{gC}}$. From here, in view of (96), we can obtain a simpler formula for the Hessian $\text{Hess}_{q,x}^{\tilde{g}}$:

$$(97) \quad \text{Hess}_{q,x}^{\tilde{g}} = \Pi_x^{\text{agC}} \circ \mathcal{D}_x \mathcal{X}_q \circ \Pi_x^{\text{elC}} : \mathcal{K}_{k,x}^{\text{agC}} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC}}.$$

Remark 5.18. The vector field $\mathcal{X}_q^{\text{gC}}$ points \tilde{g} -orthogonally to the S^1 -orbits on W_{k-1} . Hence, when x is a stationary point of $\mathcal{X}_q^{\text{gC}}$, the image of $\mathcal{D}^{\tilde{g}} \mathcal{X}_q^{\text{gC}}$ must be contained in $\mathcal{K}_{k-1}^{\text{agC}}$. Further, in this case the derivative of $\mathcal{X}_q^{\text{gC}}$ is independent of which connection we use. Therefore, when x is stationary, we can write

$$(98) \quad \text{Hess}_{q,x}^{\tilde{g}} = \mathcal{D}_x \mathcal{X}_q^{\text{gC}} = \mathcal{D}_x^{\tilde{g}} \mathcal{X}_q^{\text{gC}} : \mathcal{K}_{k,x}^{\text{agC}} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC}}.$$

It is shown in [KM07, Proposition 12.3.1] that $\text{Hess}_{q,x}$ is a Fredholm operator of index zero. Similarly, using the \tilde{g} -Hessian we have:

Lemma 5.19. *For any $x = (a, \phi) \in W_k$ with $\phi \neq 0$, the operator $\text{Hess}_{q,x}^{\tilde{g}} : \mathcal{K}_{k,x}^{\text{agC}} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC}}$ is Fredholm of index zero. Therefore, it is surjective if and only if it is injective.*

Proof. We see from (97) that $\text{Hess}_{q,x}^{\tilde{g}}$ is the conjugate of $\text{Hess}_{q,x}$ by the isomorphism $\Pi_x^{\text{agC}} : \mathcal{K}_{k,x} \rightarrow \mathcal{K}_{k,x}^{\text{agC}}$, with inverse Π_x^{elC} (see Lemma 5.5). Since $\text{Hess}_{q,x}$ is Fredholm of index 0, so is $\text{Hess}_{q,x}^{\tilde{g}}$. \square

At this point it is helpful to recall the proof of [KM07, Proposition 12.3.1], which says that $\text{Hess}_{q,x}$ is a Fredholm operator of index zero. That proof will serve as a model for some of our later arguments.

We start by recalling [KM07, Definition 12.2.1]:

Definition 5.20. *An operator L is called k -almost self-adjoint first-order elliptic (k -ASAFOE) if it is of the form*

$$L = L_0 + h$$

where

- L_0 is a first-order, self-adjoint, elliptic differential operator (with smooth coefficients) acting on sections of a vector bundle $E \rightarrow Y$;
- $h : C^\infty(Y; E) \rightarrow L^2(Y; E)$ is a linear operator on sections of E that extends to a bounded map on $L_j^2(Y; E)$ for all j with $|j| \leq k$.

For $|j| \leq k$, note that a k -ASAFOE operator $L_0 + h$ is Fredholm of index zero when viewed as a map $L_j^2(Y; E) \rightarrow L_{j-1}^2(Y; E)$. Indeed, this statement is true for L_0 (by ellipticity and self-adjointness), and adding h is just a compact perturbation.

Note that $\text{Hess}_{\mathbf{q},x}$ is not k -ASAFOE, because it does not act on *all* sections of a vector bundle. The remedy in [KM07] was to introduce an extended Hessian

$$\widehat{\text{Hess}}_{\mathbf{q},x} : \mathcal{T}_{k,x} \oplus L_k^2(Y; i\mathbb{R}) \rightarrow \mathcal{T}_{k-1,x} \oplus L_{k-1}^2(Y; i\mathbb{R})$$

by the formula

$$(99) \quad \widehat{\text{Hess}}_{\mathbf{q},x} = \begin{pmatrix} \mathcal{D}_x \mathcal{X}_{\mathbf{q}} & \mathbf{d}_x \\ \mathbf{d}_x^* & 0 \end{pmatrix}.$$

Here, \mathbf{d}_x encodes the infinitesimal gauge action at $x = (a, \phi)$:

$$\mathbf{d}_x(\xi) = (-d\xi, \xi\phi)$$

and its adjoint

$$\mathbf{d}_x^*(b, \psi) = -d^*b + i \operatorname{Re} \langle i\phi, \psi \rangle$$

can be used to define the local Coulomb slice $\mathcal{K}_{k,x}$ by the condition $\mathbf{d}_x^* = 0$.

The extended Hessian is a self-adjoint k -ASAFOE operator acting on sections of the bundle $iT^*Y \oplus \mathbb{S} \oplus i\mathbb{R}$. With respect to the decomposition $\mathcal{T}_{j,x} \oplus L_j^2(Y; i\mathbb{R}) = \mathcal{J}_{j,x} \oplus \mathcal{K}_{j,x} \oplus L_j^2(Y; i\mathbb{R})$, with $j = k$ for the domain and $j = k - 1$ for the target, we can write¹⁰

$$\widehat{\text{Hess}}_{\mathbf{q},x} = \begin{pmatrix} h_1 & h_2 & \mathbf{d}_x \\ h_3 & \text{Hess}_{\mathbf{q},x} & 0 \\ \mathbf{d}_x^* & 0 & 0 \end{pmatrix}.$$

Here, the h_i terms preserve Sobolev regularity and hence are compact as maps $L_k^2 \rightarrow L_{k-1}^2$. (As an aside, note also that the h_i terms are zero at stationary points.) Thus, dropping h_1, h_2 and h_3 from the expression above yields a Fredholm operator of index 0. In turn, since the h_i are compact, this shows that $\text{Hess}_{\mathbf{q},x}$ must be Fredholm of index 0.

5.5.2. The Hessian on the blow-up. We now define the analogous Hessian on the blow-up. In Section 12.4 of [KM07], the Hessian on the blow-up is defined by:

$$\text{Hess}_{\mathbf{q},x}^\sigma = \Pi_x^{\text{IC},\sigma} \circ \mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^\sigma : \mathcal{K}_{k,x}^\sigma \rightarrow \mathcal{K}_{k-1,x}^\sigma.$$

Note that since $\mathcal{C}_k^\sigma(Y)$ is only a submanifold of a vector space, the derivative of the section $\mathcal{X}_{\mathbf{q}}^\sigma$ need not land in its tangent bundle \mathcal{T}_{k-1}^σ . What we actually mean¹¹ by $\mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^\sigma$ is the composition of the derivative in the larger vector space with the L^2 orthogonal projection onto \mathcal{T}_{k-1}^σ . To obtain the Hessian, we further project to the corresponding local Coulomb slice by $\Pi_x^{\text{IC},\sigma}$.

To show that $\text{Hess}_{\mathbf{q},x}^\sigma$ is Fredholm of index zero, Kronheimer and Mrowka introduced the extended Hessian in the blow-up:

$$(100) \quad \widehat{\text{Hess}}_{\mathbf{q},x}^\sigma = \begin{pmatrix} \mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^\sigma & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\sigma,\dagger} & 0 \end{pmatrix}.$$

¹⁰In [KM07, Equation (12.7)], the term h_1 was inadvertently missing.

¹¹In [KM07], the notation \mathcal{D} is used instead of \mathcal{D}^σ . We use \mathcal{D}^σ to distinguish this derivative from the L^2 derivative \mathcal{D} in the ambient vector space $L^2(Y; iT^*Y) \oplus \mathbb{R} \oplus L^2(Y; \mathbb{S})$, which we will also work with.

Here,

$$(101) \quad \mathbf{d}_x^\sigma(\xi) = (-d\xi, 0, \xi\phi),$$

$$(102) \quad \mathbf{d}_x^{\sigma,\dagger}(b, r, \psi) = -d^*b + is^2 \operatorname{Re}\langle i\phi, \psi \rangle + i|\phi|^2 \operatorname{Re} \mu_Y(\langle i\phi, \psi \rangle),$$

so that $\mathcal{K}_j^\sigma = \ker \mathbf{d}_x^{\sigma,\dagger}$, $\mathcal{J}_j^\sigma = \operatorname{im} \mathbf{d}_x^\sigma$ are the local Coulomb slice and the tangent to the gauge orbit, as in (45). Note that on the blow-up, $\mathbf{d}_x^{\sigma,\dagger}$ is not quite the adjoint to \mathbf{d}_x^σ , just as the decomposition $\mathcal{T}_j^\sigma = \mathcal{J}_j^\sigma \oplus \mathcal{K}_j^\sigma$ is not L^2 orthogonal. Therefore, unlike the extended Hessian $\widehat{\operatorname{Hess}}_{q,x}$, the operator $\widehat{\operatorname{Hess}}_{q,x}^\sigma$ is not symmetric. Moreover, $\widehat{\operatorname{Hess}}_{q,x}^\sigma$ does not a priori act on sections of a vector bundle, but rather on the subspace of $\mathcal{T}_{j,x}^\sigma \oplus L_j^2(Y; i\mathbb{R})$, where we have the condition $\operatorname{Re}\langle \phi, \psi \rangle_{L^2} = 0$. Nevertheless, we can combine the r and ψ components of $(b, r, \psi) \in \mathcal{T}_{j,x}^\sigma$ into

$$\psi = \psi + r\phi,$$

and then we can think of $\widehat{\operatorname{Hess}}_{q,x}^\sigma$ as acting on sections of $iT^*Y \oplus \mathbb{S} \oplus i\mathbb{R}$. It then becomes a k -ASAFOE operator (hence Fredholm of index zero), and starting from here one can show that $\operatorname{Hess}_{q,x}^\sigma$ is Fredholm of index zero—much as in the proof of Lemma 5.19. Further, when x is a non-degenerate stationary point, $\operatorname{Hess}_{q,x}^\sigma$ is invertible and has real spectrum (even though it is not self-adjoint). See [KM07, Section 12.4] for details.

Let us now move to global Coulomb gauge. Fix $x = (a, s, \phi) \in W^\sigma$. In this setting, we define a \tilde{g} -Hessian in the blow-up by:

$$(103) \quad \operatorname{Hess}_{q,x}^{\tilde{g},\sigma} = \Pi_x^{\operatorname{agC},\sigma} \circ \mathcal{D}_x^{\tilde{g},\sigma} \mathcal{X}_q^{\operatorname{gC},\sigma} : \mathcal{K}_{k,x}^{\operatorname{agC},\sigma} \rightarrow \mathcal{K}_{k-1,x}^{\operatorname{agC},\sigma},$$

where

$$(104) \quad \mathcal{D}^{\tilde{g},\sigma} X := \Pi_*^{\operatorname{gC},\sigma} \circ \mathcal{D}^\sigma(\Pi^{\operatorname{elC},\sigma}(X)) \circ \Pi^{\operatorname{elC},\sigma}.$$

is the connection on $\mathcal{T}^{\operatorname{gC},\sigma}$ defined by analogy with (96).

Since $\mathcal{X}_q^{\operatorname{gC},\sigma} = \Pi_*^{\operatorname{gC},\sigma} \circ \mathcal{X}_q^\sigma$, in view of Lemma 5.6(a) we have $\mathcal{X}_q^\sigma = \Pi^{\operatorname{elC},\sigma} \circ \mathcal{X}_q^{\operatorname{gC},\sigma}$. Therefore, we can re-write the blown-up \tilde{g} -Hessian as

$$(105) \quad \operatorname{Hess}_{q,x}^{\tilde{g},\sigma} = \Pi_x^{\operatorname{agC},\sigma} \circ \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma \circ \Pi_x^{\operatorname{elC},\sigma}.$$

Remark 5.21. When x is on the blow-up locus (i.e., $s = 0$), we know from Lemma 5.5 (b) that the anticircular global Coulomb slice coincides with the local Coulomb slice and $\Pi_x^{\operatorname{elC},\sigma} = \Pi_x^{\operatorname{agC},\sigma}$. We see from (105) that $\operatorname{Hess}_{q,x}^{\tilde{g},\sigma}$ agrees with the ordinary blow-up Hessian $\operatorname{Hess}_{q,x}^\sigma$.

Lemma 5.22. (a) For any $x \in W_k^\sigma$, the operator $\operatorname{Hess}_{q,x}^{\tilde{g},\sigma}$ is Fredholm of index zero.

(b) When x is a non-degenerate stationary point of $\mathcal{X}_q^{\operatorname{gC},\sigma}$, the operator $\operatorname{Hess}_{q,x}^{\tilde{g},\sigma}$ is invertible and has real spectrum.

Proof. When x is reducible, we have $\operatorname{Hess}_{q,x}^{\tilde{g},\sigma} = \operatorname{Hess}_{q,x}^\sigma$, and the corresponding results for $\operatorname{Hess}_{q,x}^\sigma$ were established in [KM07, Section 12.4].

When x is irreducible, $\operatorname{Hess}_{q,x}^{\tilde{g},\sigma}$ is conjugate to $\operatorname{Hess}_{q,x}^{\tilde{g}}$ via the blow-down map. Thus, part (a) is a consequence of Lemma 5.19. For part (b), since x is stationary, by (97), we have that $\operatorname{Hess}_{q,x}^{\tilde{g}} = \mathcal{D}_x^{\tilde{g}} \mathcal{X}_q^{\operatorname{gC}}$, and the latter operator is invertible and self-adjoint (being a formal Hessian). The conclusion follows. \square

Let us also mention:

Lemma 5.23. *The map $\text{Hess}_{\mathbf{q}}^{\tilde{g},\sigma} : \mathcal{K}_k^{\text{agC},\sigma} \rightarrow \mathcal{K}_{k-1}^{\text{agC},\sigma}$ is a continuous bundle map.*

Proof. This follows from the continuity of the factors in (105). In particular, recall that $\Pi^{\text{agC},\sigma}$ is the composition of $\Pi_*^{\text{gC},\sigma}$ with the projection (74) onto $\mathcal{K}^{\text{agC},\sigma}$. The continuity of $\Pi_*^{\text{gC},\sigma}$ can be proved by an argument similar to that in Lemma 3.1. As for the projection (74), it can be seen from that formula that it is a continuous operation. Finally, $\Pi^{\text{elC},\sigma}$ is continuous by the same arguments as in Lemma 3.3. \square

5.5.3. *The split extended Hessian on the blow-up.* For the purposes of this paper, it is helpful to work with a different extended Hessian than the one in (100). Precisely, we replace (102) with an operator

$$\mathbf{d}_x^{\text{sp},\sigma,\dagger} : \mathcal{T}_{j,x}^\sigma \rightarrow L_j^2(Y; i\mathbb{R}),$$

defined as follows. We decompose the domain $\mathcal{T}_{j,x}^\sigma = \mathcal{K}_{j,x}^\sigma \oplus \mathcal{J}_{j,x}^\sigma$ into three summands $\mathcal{K}_{j,x}^\sigma \oplus \mathcal{J}_{j,x}^{\text{gC},\sigma} \oplus \mathcal{J}_{j,x}^{\circ,\sigma}$, and also decompose the codomain $L_{j-1}^2(Y; i\mathbb{R})$ into $(\text{im } d^*)_{j-1} \oplus i\mathbb{R}$ (that is, into functions that integrate to zero and constant functions). With respect to these decompositions, we let

$$(106) \quad \mathbf{d}_x^{\text{sp},\sigma,\dagger} = \begin{pmatrix} 0 & 0 & -d^* \\ 0 & i \text{Re}\langle i\phi, \cdot \rangle_{L^2} & 0 \end{pmatrix},$$

where $-d^*$ acts on the first component of $(-d\xi, 0, \xi\phi) \in \mathcal{J}_{j,x}^{\circ,\sigma}$, and $i \text{Re}\langle i\phi, \cdot \rangle_{L^2}$ acts on the last component of $(0, 0, it\phi) \in \mathcal{J}_{j,x}^{\text{gC},\sigma}$, $t \in \mathbb{R}$. Note that $\|\phi\|_{L^2} = 1$, so $i \text{Re}\langle i\phi, it\phi \rangle_{L^2}$ simply equals it .

One can also write a more compressed formula for $\mathbf{d}_x^{\text{sp},\sigma,\dagger}$. Recall from (66) that $\Pi_{(a,s,\phi)}^{\text{IC},\sigma}(b, r, \psi) = (b - d\zeta, r, \psi + \zeta\phi)$, for some $\zeta = \zeta(x, b, r, \psi) : Y \rightarrow i\mathbb{R}$. Then:

$$\mathbf{d}_x^{\text{sp},\sigma,\dagger}(b, r, \psi) = d^*d\zeta + \mu_Y(\zeta).$$

Note that the kernel of $\mathbf{d}_x^{\text{sp},\sigma,\dagger}$ is the local Coulomb slice $\mathcal{K}_{j,x}^\sigma$, just as for $\mathbf{d}_x^{\sigma,\dagger}$.

We now define the *split extended Hessian in the blow-up* to be

$$(107) \quad \widehat{\text{Hess}}_{\mathbf{q},x}^{\text{sp},\sigma} = \begin{pmatrix} \mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^\sigma & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\text{sp},\sigma,\dagger} & 0 \end{pmatrix} : \mathcal{T}_{j,x}^\sigma \oplus L_j^2(Y; i\mathbb{R}) \rightarrow \mathcal{T}_{j-1,x}^\sigma \oplus L_{j-1}^2(Y; i\mathbb{R}).$$

At any x , by combining ψ and r into $\boldsymbol{\psi} = \psi + r\phi$ as in the case of $\widehat{\text{Hess}}_{\mathbf{q},x}^\sigma$, we see that $\widehat{\text{Hess}}_{\mathbf{q},x}^{\text{sp},\sigma}$ is a k -ASAFOE operator. We also have:

Lemma 5.24. *If x is a non-degenerate stationary point of $\mathcal{X}_{\mathbf{q}}^{\text{gC},\sigma}$, then $\widehat{\text{Hess}}_{\mathbf{q},x}^{\text{sp},\sigma}$ is invertible and has real spectrum.*

Proof. The argument is similar to that in [KM07, proof of Lemma 12.4.3]. The operator $\widehat{\text{Hess}}_{\mathbf{q},x}^{\text{sp},\sigma}$ has a block form where one block is $\text{Hess}_{\mathbf{q},x}^\sigma$ and the other is

$$(108) \quad \begin{pmatrix} 0 & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\text{sp},\sigma,\dagger} & 0 \end{pmatrix}.$$

It is established in [KM07, proof of Lemma 12.4.3] that the operator $\text{Hess}_{\mathbf{q},x}^\sigma$ is invertible and has real spectrum. To justify that (108) is invertible and has real spectrum, it suffices to prove that the operator $\mathbf{d}_x^{\text{sp},\sigma,\dagger} \mathbf{d}_x^\sigma$ is self-adjoint and strictly positive.

To see this last fact, we compute

$$\mathbf{d}_x^{\text{sp},\sigma,\dagger} \mathbf{d}_x^\sigma(\xi) = \mathbf{d}_x^{\text{sp},\sigma,\dagger}(-d\xi, 0, \xi\phi) = \Delta\xi + \mu_Y(\xi).$$

With respect to the orthogonal decomposition $L_{j-1}^2(Y; i\mathbb{R}) = (\text{im } d^*)_{j-1} \oplus i\mathbb{R}$, we have

$$\mathbf{d}_x^{\text{sp}, \sigma, \dagger} \mathbf{d}_x^\sigma = \begin{pmatrix} \Delta & 0 \\ 0 & 1 \end{pmatrix}.$$

Both diagonal entries, Δ and 1, are self-adjoint and positive on the respective summands. This completes the proof. \square

Remark 5.25. We could have defined a split extended Hessian in the blow-down as well (at irreducibles), but we had no need for that construction.

5.5.4. The \tilde{g} -extended Hessian on the blow-up. We now construct another extended Hessian in the blow-up, using the \tilde{g} metric. The definition is somewhat similar to that of the split extended Hessian in Section 5.5.3 above. Precisely, we set

$$(109) \quad \widehat{\text{Hess}}_{\mathbf{q}, x}^{\tilde{g}, \sigma} = \begin{pmatrix} S_x \circ \mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^\sigma \circ S_x^{-1} & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\sigma, \tilde{\dagger}} & 0 \end{pmatrix}.$$

Here, S_x is the shear map from (79), and to define $\mathbf{d}_x^{\sigma, \tilde{\dagger}} : \mathcal{T}_{j,x}^\sigma \rightarrow L_{j-1}^2(Y; i\mathbb{R})$, we use the decompositions

$$\mathcal{T}_{j,x}^\sigma = \mathcal{K}_{j,x}^{\text{agC}, \sigma} \oplus \mathcal{J}_{j,x}^{\text{gC}, \sigma} \oplus \mathcal{J}_{j,x}^{\circ, \sigma}, \quad L_{j-1}^2(Y; i\mathbb{R}) = (\text{im } d^*)_{j-1} \oplus i\mathbb{R},$$

and set

$$(110) \quad \mathbf{d}_x^{\sigma, \tilde{\dagger}} = \begin{pmatrix} 0 & 0 & -d^* \\ 0 & i\langle i\phi, \cdot \rangle_{\tilde{g}} & 0 \end{pmatrix}.$$

Alternatively, we have the formula

$$(111) \quad \mathbf{d}_x^{\sigma, \tilde{\dagger}}(b, r, \psi) = -d^*b + i\langle i\phi, \psi + (Gd^*b)\phi \rangle_{\tilde{g}}.$$

Observe the kernel of $\mathbf{d}_x^{\sigma, \tilde{\dagger}}$ is the anticircular global Coulomb slice $\mathcal{K}_{j,x}^{\text{agC}, \sigma}$, given by the conditions $d^*b = 0$ and $\langle i\phi, \psi \rangle_{\tilde{g}} = 0$.

At any x , by combining ψ and r into $\psi = \psi + r\phi$ as before, we would like to claim that $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\tilde{g}, \sigma}$ is a k -ASAFOE operator as a first step to proving that $\widehat{\text{Hess}}_{\mathbf{q}, x}^\sigma$ is invertible and has real spectrum at non-degenerate stationary points. It turns out that if x is not a stationary point, then $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\tilde{g}, \sigma}$ is only $(k-1)$ -ASAFOE, so we will have to be careful. To establish these properties, we will use a different operator which agrees with $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\tilde{g}, \sigma}$ at stationary points. This operator will also be useful in Section 5.9 and is described in the following lemma.

Lemma 5.26. *Let $x \in W_k^\sigma$ and $1 \leq j \leq k$. Consider the operator $\mathcal{H}_x^\sigma : \mathcal{T}_{j,x}^\sigma \oplus L_j^2(Y; i\mathbb{R}) \rightarrow \mathcal{T}_{j-1,x}^\sigma \oplus L_{j-1}^2(Y; i\mathbb{R})$ given in block form by*

$$\mathcal{H}_x^\sigma = \begin{pmatrix} (\mathcal{D}_x^\sigma \mathcal{X}_{\mathbf{q}}^{\text{gC}, \sigma}) \circ \Pi_*^{\text{gC}, \sigma} & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\sigma, \tilde{\dagger}} & 0 \end{pmatrix}.$$

(a) *Under the identification of $\mathcal{T}_{j,x}^\sigma \oplus L_j^2(Y; i\mathbb{R})$ with $L_j^2(Y; iT^*Y \oplus \mathbb{S} \oplus \mathbb{R})$ and likewise for $j-1$, the operator \mathcal{H}_x^σ is $(k-1)$ -ASAFOE with linear part*

$$L_0 = \begin{pmatrix} *d & 0 & -d \\ 0 & D & 0 \\ -d^* & 0 & 0 \end{pmatrix};$$

- (b) when $j = k$, the operator \mathcal{H}_x^σ differs from L_0 by a compact operator from L_k^2 to L_{k-1}^2 ;
- (c) if x is a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$, then \mathcal{H}_x^σ is k -ASAFOE;
- (d) if x is a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$, then $\mathcal{H}_x^\sigma = \widehat{\text{Hess}}_{q,x}^{\text{g},\sigma}$.

Proof. (a) For notation, we write $L^\sigma : \mathcal{T}_{j,x}^\sigma \oplus L_j^2(Y; i\mathbb{R}) \rightarrow \mathcal{T}_{j-1,x}^\sigma \oplus L_{j-1}^2(Y; i\mathbb{R})$ for the operator induced by L_0 . Our goal is to show that \mathcal{H}_x^σ differs from L^σ by bounded operators from L_j^2 to L_j^2 for $1 \leq j \leq k-1$. (Technically, we must show that these are induced by operators from C^∞ to L^2 , but this will be clear from the explicit description.) We break up the analysis of \mathcal{H}_x^σ into how it acts on $L_j^2(Y; i\mathbb{R})$, $\mathcal{T}_{j,x}^{\text{gC},\sigma}$ and $\mathcal{J}_{j,x}^{\circ,\sigma}$. In fact, we will show the difference is bounded from L_j^2 to L_j^2 even when $j = k$, except for one term.

First, consider $\beta \in L_j^2(Y; i\mathbb{R})$. We have that $\mathcal{H}_x^\sigma(\beta) = (-d\beta, 0, \beta\phi)$. Then, we have $\mathcal{H}_x^\sigma(\beta) - L^\sigma(\beta) = \beta\phi$, which is bounded as a linear map from L_j^2 to L_j^2 by Sobolev multiplication whenever $1 \leq j \leq k$ since $\phi \in L_k^2$.

Next, let $v = (-d\xi, 0, \xi\phi) \in \mathcal{J}_{j,x}^{\circ,\sigma}$. Then, since v is in the kernel of the infinitesimal global Coulomb projection $\Pi_*^{\text{gC},\sigma}$, we have

$$\mathcal{H}_x^\sigma(v) = (0, \mathbf{d}_x^{\sigma,\dagger}(v)) = (0, d^*d\xi) \in \mathcal{T}_{j-1,x}^\sigma \oplus L_{j-1}^2(Y; i\mathbb{R}).$$

In the second component, $\mathcal{H}_x^\sigma(v)$ and $L^\sigma(v)$ agree. Thus, it remains to show that the component of $L^\sigma(v)$ landing in $\mathcal{T}_{j-1,x}^\sigma$ is bounded from L_j^2 to L_j^2 for $1 \leq j \leq k-1$. We have

$$L^\sigma(v) - \mathcal{H}_x^\sigma(v) = (0, \text{Re}\langle D(\xi\phi), \phi \rangle_{L^2}, D(\xi\phi) - \text{Re}\langle D(\xi\phi), \phi \rangle_{L^2}\phi) \in \mathcal{T}_{j-1,x}^\sigma.$$

This differs from $v \mapsto (0, 0, D(\xi\phi))$ by a bounded operator from L_j^2 to L_k^2 , so we focus on $D(\xi\phi) = \rho(d\xi)\phi + \xi D\phi$. Since $v \in L_j^2$, we see that $d\xi \in L_j^2$ and thus $\rho(d\xi)\phi$ is bounded in L_j^2 by Sobolev multiplication. Therefore, $v \mapsto (0, 0, \rho(d\xi)\phi)$ is bounded from L_j^2 to L_j^2 (even if $j = k$). For the term $\xi D\phi$, we note that $D\phi \in L_{k-1}^2$, so the map $(-d\xi, 0, \xi\phi) \mapsto \xi D\phi$ is bounded as a linear map from L_j^2 to L_j^2 as long as $j \leq k-1$. This establishes the desired form for $v \in \mathcal{J}_{j,x}^{\circ,\sigma}$.

It thus remains to compare \mathcal{H}_x^σ and L^σ on $\mathcal{T}_{j,x}^{\text{gC},\sigma}$. Let $v = (b, r, \psi) \in \mathcal{T}_{j,x}^{\text{gC},\sigma}$. First, note that the component of $\mathcal{H}_x^\sigma(v)$ landing in $L_{j-1}^2(Y; i\mathbb{R})$ is given by $i\langle i\phi, \psi \rangle_{\bar{g}}$ which is bounded as a map from $\mathcal{T}_{j,x}^{\text{gC},\sigma}$ to $L_j^2(Y; i\mathbb{R})$ for $j \leq k$; this is compatible with the fact that $L^\sigma(b, r, \psi)$ has no component landing in $L_{j-1}^2(Y; i\mathbb{R})$ since $d^*b = 0$. Thus, it remains to focus on the component of $\mathcal{H}_x^\sigma(v)$ contained in $\mathcal{T}_{j,x}^\sigma$. Since $\widehat{\text{Hess}}_{q,x}^\sigma$ is k -ASAFOE with linear term also given by L_0 , it suffices to show that $\mathcal{D}_x^\sigma \mathcal{X}_q^\sigma$ differs from $\mathcal{D}_x^\sigma \mathcal{X}_q^{\text{gC},\sigma}$ by bounded operators from L_j^2 to L_j^2 . Direct computation shows

$$\mathcal{X}_q^\sigma(a, s, \phi) - \mathcal{X}_q^{\text{gC},\sigma}(a, s, \phi) = (dGd^*\mathbf{q}^0(a, s\phi), 0, -(Gd^*\mathbf{q}^0(a, s\phi))\phi).$$

Recall that to compute \mathcal{D}^σ , we compute the L^2 derivative as a map into the affine space $L^2(Y; iT^*Y) \oplus \mathbb{R} \oplus L^2(Y; \mathbb{S})$ and then apply L^2 projection to $\mathcal{T}_{0,x}^\sigma$. We first compute the affine derivative

$$\begin{aligned} (\mathcal{D}_{(a,s,\phi)} \mathcal{X}_q^\sigma - \mathcal{D}_{(a,s,\phi)} \mathcal{X}_q^{\text{gC},\sigma})(b, r, \psi) &= (dGd^*\mathcal{D}_{(a,s,\phi)} \mathbf{q}^0(b, r\phi + s\psi), 0, \\ &\quad - Gd^*\mathcal{D}_{(a,s,\phi)} \mathbf{q}^0(b, r\phi + s\psi)\phi - Gd^*\mathbf{q}^0(a, s\phi)\psi). \end{aligned}$$

Projecting, we obtain

(112)

$$\begin{aligned} \left(\mathcal{D}_{(a,s,\phi)}^\sigma \mathcal{X}_q^\sigma - \mathcal{D}_{(a,s,\phi)}^\sigma \mathcal{X}_q^{\text{gC},\sigma} \right) (b, r, \psi) = & (dGd^* \mathcal{D}_{(a,s\phi 0)} \mathbf{q}^0(b, r\phi + s\psi), 0, \\ & - Gd^* \mathcal{D}_{(a,s\phi 0)} \mathbf{q}^0(b, r\phi + s\psi)\phi - Gd^* \mathbf{q}^0(a, s\phi)\psi + \\ & \text{Re}\langle Gd^* \mathbf{q}^0(a, s\phi)\psi, \phi \rangle_{L^2\phi}). \end{aligned}$$

Here we are using that $Gd^* \mathcal{D}_{(a,s\phi)} \mathbf{q}^0(b, r\phi + s\psi)\phi$ is real L^2 orthogonal to ϕ since \mathbf{q}^0 is purely imaginary. This operator is seen to be bounded from L_j^2 to L_j^2 for $j \leq k$ since \mathbf{q} is tame and because G raises Sobolev regularity by 2.

(b) From the above argument, we saw that \mathcal{H}_x^σ differs from L^σ by a bounded map from L_j^2 to L_j^2 for all $1 \leq j \leq k$, except possibly for a single term: $(-d\xi, 0, \xi\phi) \mapsto \xi D\phi$ where $(-d\xi, 0, \xi\phi) \in \mathcal{J}_{k,x}^{\circ,\sigma}$. By postcomposing with the compact inclusion of L_j^2 into L_{j-1}^2 , we have the desired result, except for this exceptional term. In this case, since $D\phi \in L_{k-1}^2$, we have that this operator is bounded as a linear map from L_{k-1}^2 to L_{k-1}^2 . Therefore, we precompose with the compact inclusion from L_k^2 to L_{k-1}^2 to obtain the desired compactness.

(c) As in the above case, we previously showed that \mathcal{H}_x^σ differs from L^σ by a bounded map from L_j^2 to L_j^2 for all $1 \leq j \leq k$, except for the term involving $D\phi$. If x is a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$, then we know that ϕ is actually in L_{k+1}^2 and thus $D\phi$ is contained in L_k^2 , and the argument proceeds with $j = k$.

(d) Let x be a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$ (and hence also \mathcal{X}_q^σ). By (109), it suffices to show that $S_x \circ \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma \circ S_x^{-1}$ agrees with $\mathcal{D}_x^\sigma \mathcal{X}_q^{\text{gC},\sigma} \circ \Pi_*^{\text{gC},\sigma}$ on $\mathcal{T}_{j,x}^\sigma$. First, suppose $v \in \mathcal{J}_{j,x}^{\circ,\sigma}$. Then, we have that $S_x^{-1}(v) = v$. It then follows that $S_x \circ \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma \circ S_x^{-1}(v) = 0$, since v is tangent to the gauge orbit at a stationary point. This agrees with the fact that the kernel of $\Pi_*^{\text{gC},\sigma}$ is $\mathcal{J}^{\circ,\sigma}$.

Next, consider the case $v \in \mathcal{T}_{j,x}^{\text{gC},\sigma}$. Since $S_x^{-1}(v) - v$ is contained in $\mathcal{J}_{j,x}^{\circ,\sigma}$, we have that $\mathcal{D}_x^\sigma \mathcal{X}_q^\sigma \circ S_x^{-1}(v) = \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v)$. Further, since x is a stationary point, we have that $\mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v) \in \mathcal{K}_{j,x}^\sigma$, and therefore $S_x(\mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v)) = \Pi_*^{\text{gC},\sigma} \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v)$. Thus, it suffices to prove that

$$(113) \quad \Pi_*^{\text{gC},\sigma} \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v) = \mathcal{D}_x^\sigma \mathcal{X}_q^{\text{gC},\sigma}(v).$$

Since $\mathcal{D}_x^\sigma \mathcal{X}_q^{\text{gC},\sigma}(v) \in \mathcal{T}_{j,x}^{\text{gC},\sigma}$, we have that (113) is equivalent to

$$(114) \quad \Pi_*^{\text{gC},\sigma} \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma(v) = \Pi_*^{\text{gC},\sigma} \mathcal{D}_x^\sigma \mathcal{X}_q^{\text{gC},\sigma}(v).$$

Thus, we can establish (114) if applying infinitesimal global Coulomb projection to (112) vanishes. We split (112) into two parts. The first term is

$$(0, 0, Gd^* \mathbf{q}^0(a, s\phi)\psi + \text{Re}\langle Gd^* \mathbf{q}^0(a, s\phi)\psi, \phi \rangle_{L^2\phi}).$$

Since x is a stationary point of \mathcal{X}_q^σ , we have that $-\mathbf{q}^0(a, s\phi) = *da$. Therefore, $Gd^* \mathbf{q}^0(a, s\phi) = 0$ and the above expression vanishes. Thus, it suffices to show that infinitesimal global Coulomb projection vanishes on the remaining term

$$(dGd^* \mathcal{D}_{(a,s\phi 0)} \mathbf{q}^0(b, r\phi + s\psi), 0, -Gd^* \mathcal{D}_{(a,s\phi 0)} \mathbf{q}^0(b, r\phi + s\psi)\phi).$$

This term is contained in $\mathcal{J}_{j,x}^\circ$, which is precisely the kernel of the infinitesimal global Coulomb projection. This completes the proof. \square

With the above lemma, we can now show that $\widehat{\text{Hess}}_{q,x}^{\tilde{g},\sigma}$ shares two important properties with $\widehat{\text{Hess}}_{q,x}^\sigma$ and $\widehat{\text{Hess}}_{q,x}^{\text{sp},\sigma}$.

Lemma 5.27. *If x is a non-degenerate stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$, then $\widehat{\text{Hess}}_{q,x}^{\tilde{g},\sigma}$ is invertible and has real spectrum.*

Proof. At stationary points, with respect to the decomposition $\mathcal{T}_{j,x}^\sigma = \mathcal{K}_{j,x}^\sigma \oplus \mathcal{J}_{j,x}^\sigma$, the derivative $\mathcal{D}_x \mathcal{X}_q^\sigma$ takes the form

$$\begin{pmatrix} \text{Hess}_{q,x}^\sigma & 0 \\ 0 & 0 \end{pmatrix}.$$

Recall from Remark 5.7 that S_x maps $\mathcal{K}_{j,x}^\sigma$ to $\mathcal{K}_{j,x}^{\text{agC},\sigma}$ and preserves $\mathcal{J}_{j,x}^\sigma$. Therefore, after conjugating by the shear S_x , we have

$$S_x \circ \mathcal{D}_x \mathcal{X}_q^\sigma \circ S_x^{-1} = \begin{pmatrix} \text{Hess}_{q,x}^{\tilde{g},\sigma} & 0 \\ 0 & 0 \end{pmatrix},$$

with respect to the decomposition $\mathcal{K}_{j,x}^{\text{agC},\sigma} \oplus \mathcal{J}_{j,x}^\sigma$.

We deduce that the operator $\widehat{\text{Hess}}_{q,x}^{\tilde{g},\sigma}$ has a block form where one block is $\text{Hess}_{q,x}^{\tilde{g},\sigma}$ and the other is

$$(115) \quad \begin{pmatrix} 0 & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\sigma,\dagger} & 0 \end{pmatrix}.$$

By Lemma 5.22, $\text{Hess}_{q,x}^{\tilde{g},\sigma}$ is invertible and has real spectrum. It remains to check that (115) is invertible and has real spectrum. As in the proof of Lemma 5.24, we do this by checking that $\mathbf{d}_x^{\sigma,\dagger} \mathbf{d}_x^\sigma$ is self-adjoint and strictly positive.

Indeed, we have

$$\mathbf{d}_x^{\sigma,\dagger} \mathbf{d}_x^\sigma(\xi) = \Delta \xi + \mu_Y(\xi) \|i\phi\|_g^2.$$

In block form with respect to the decomposition $L_{j-1}^2(Y; i\mathbb{R}) = (\text{im } d^*)_{j-1} \oplus i\mathbb{R}$, we have

$$\mathbf{d}_x^{\text{sp},\sigma,\dagger} \mathbf{d}_x^\sigma = \begin{pmatrix} \Delta & 0 \\ 0 & \|i\phi\|_g^2 \end{pmatrix}.$$

Both diagonal entries are self-adjoint and positive. (Note that ϕ is nonzero, because it was normalized to have unit L^2 norm.) \square

5.5.5. Interpolations. In the proof of Proposition 5.40 below, we will need to interpolate between the extended Hessians $\widehat{\text{Hess}}_{q,x}^\sigma$ and $\widehat{\text{Hess}}_{q,x}^{\tilde{g},\sigma}$, by going through k -ASAFOE operators that are still invertible and have real spectrum.

We can do this in two steps. First, we interpolate linearly between $\widehat{\text{Hess}}_{q,x}^\sigma$ and the split extended Hessian $\widehat{\text{Hess}}_{q,x}^{\text{sp},\sigma}$.

Lemma 5.28. *If x is a non-degenerate stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$, then for any $\rho \in [0, 1]$, we have that*

$$(1 - \rho) \cdot \widehat{\text{Hess}}_{q,x}^\sigma + \rho \cdot \widehat{\text{Hess}}_{q,x}^{\text{sp},\sigma}$$

is invertible and has real spectrum.

Proof. We use a block decomposition as in the proofs of Lemma 12.4.3 in [KM07] and of Lemma 5.24 above. It suffices to check that

$$(1 - \rho) \cdot \mathbf{d}_x^{\sigma, \dagger} \mathbf{d}_x^\sigma + \rho \cdot \mathbf{d}_x^{\text{sp}, \sigma, \dagger} \mathbf{d}_x^\sigma$$

is self-adjoint and strictly positive. This is true because both terms are self-adjoint and strictly positive. \square

For the second step, we interpolate between $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\text{sp}, \sigma}$ and $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\tilde{g}, \sigma}$. We do this by considering the family of metrics on $\mathcal{T}_{j, x}$ given by

$$g_\rho = (1 - \rho) \cdot g_{L^2} + \rho \cdot \tilde{g}, \quad \rho \in [0, 1],$$

where g_{L^2} denotes the L^2 metric. We consider the g_ρ -orthogonal complements to $\mathcal{J}_{j, x}$ and $\mathcal{J}_{j, x}^\circ$, which we denote by $\mathcal{K}_{j, x}^\rho$ and $\mathcal{K}_{j, x}^{\rho, \text{e}}$, respectively. After blowing-up, we obtain a complement $\mathcal{K}_{j, x}^{\rho, \sigma}$ to $\mathcal{J}_{j, x}^\sigma$ and a complement $\mathcal{K}_{j, x}^{\rho, \text{e}, \sigma}$ to $\mathcal{J}_{j, x}^{\circ, \sigma}$. We construct the shear map S_x^ρ that takes $\mathcal{K}_{j, x}^{\text{e}, \sigma}$ into $\mathcal{K}_{j, x}^{\rho, \text{e}, \sigma}$ and is the identity on $\mathcal{J}_{j, x}^{\circ, \sigma}$. Further, we let

$$(116) \quad \mathbf{d}_x^{\rho, \sigma, \dagger} = \begin{pmatrix} 0 & 0 & -d^* \\ 0 & i\langle i\phi, \cdot \rangle_{g_\rho} & 0 \end{pmatrix} : \mathcal{K}_{j, x}^{\rho, \sigma} \oplus \mathcal{J}_{j, x}^{\text{gC}, \sigma} \oplus \mathcal{J}_{j, x}^{\circ, \sigma} \rightarrow (\text{im } d^*)_{j-1} \oplus i\mathbb{R},$$

so that $\mathbf{d}_x^{0, \sigma, \dagger} = \mathbf{d}_x^{\text{sp}, \sigma, \dagger}$ and $\mathbf{d}_x^{1, \sigma, \dagger} = \mathbf{d}_x^{\sigma, \dagger}$. Finally, define

$$(117) \quad \widehat{\text{Hess}}_{\mathbf{q}, x}^{\rho, \sigma} = \begin{pmatrix} S_x^\rho \circ \mathcal{D}_x^\sigma \mathcal{X}_q^\sigma \circ (S_x^\rho)^{-1} & \mathbf{d}_x^\sigma \\ \mathbf{d}_x^{\rho, \sigma, \dagger} & 0 \end{pmatrix}.$$

The same arguments as in Lemmas 5.27 and 5.28 give the following:

Lemma 5.29. *If x is a non-degenerate stationary point of $\mathcal{X}_q^{\text{gC}, \sigma}$, and $\rho \in [0, 1]$, then $\widehat{\text{Hess}}_{\mathbf{q}, x}^{\rho, \sigma}$ is invertible and has real spectrum.*

5.6. Non-degeneracy of stationary points in Coulomb gauge. Recall that an irreducible stationary point $x = (a, \phi)$ of \mathcal{X}_q is non-degenerate if \mathcal{X}_q is transverse to the gauge orbit at x or, equivalently, transverse to the subbundle \mathcal{J}_{k-1} at x . We would like to rephrase this condition both in terms of Coulomb gauge and in terms of Hessians.

Lemma 5.30. *Let $x \in W_k$ be an irreducible stationary point of $\mathcal{X}_q^{\text{gC}}$. The following are equivalent:*

- (i) x is non-degenerate (i.e. \mathcal{X}_q is transverse to \mathcal{J}_{k-1} at x),
- (ii) $\text{Hess}_{\mathbf{q}, x} : \mathcal{K}_{k, x} \rightarrow \mathcal{K}_{k-1, x}$ is surjective,
- (iii) $\mathcal{X}_q^{\text{gC}}$ is transverse to $\mathcal{J}_x^{\text{gC}}$,
- (iv) $\text{Hess}_{\mathbf{q}, x}^{\tilde{g}} : \mathcal{K}_{k, x}^{\text{agC}} \rightarrow \mathcal{K}_{k-1, x}^{\text{agC}}$ is surjective.

Proof. The equivalence (i) \iff (ii) is proved in [KM07, Lemma 12.4.1]. It is a consequence of the fact that at a critical point, with respect to the decompositions $\mathcal{T}_j = \mathcal{J}_j \oplus \mathcal{K}_j$ (with $j = k$ for the domain and $j = k - 1$ for the image), the derivative $\mathcal{D}_x \mathcal{X}_q$ has the block form

$$\begin{pmatrix} 0 & 0 \\ 0 & \text{Hess}_{\mathbf{q}, x} \end{pmatrix}.$$

For the equivalence (iii) \iff (iv), we apply a similar reasoning: $\mathcal{D}_x^{\tilde{g}} \mathcal{X}_q^{\text{gC}}$ vanishes on $\mathcal{J}_x^{\text{gC}}$, the tangents to the S^1 -orbits in W_k ; therefore, $\mathcal{D}_x^{\tilde{g}} \mathcal{X}_q^{\text{gC}}$ has a block form with respect

to the decomposition $\mathcal{T}_j^{\text{gC}} = \mathcal{J}^{\text{gC}} \oplus \mathcal{K}_j^{\text{gC}}$ from (60); and the only nonzero entry in this block form is $\text{Hess}_{\mathfrak{q},x}^{\tilde{g}}$. (Since we are at a stationary point, note that in fact $\mathcal{D}_x^{\tilde{g}} \mathcal{X}_{\mathfrak{q}}^{\text{gC}} = \mathcal{D}_x \mathcal{X}_{\mathfrak{q}}^{\text{gC}}$.)

Finally, for the equivalence (ii) \iff (iv), recall from (97) and the proof of Lemma 5.19 that $\text{Hess}_{\mathfrak{q},x}^{\tilde{g}}$ is the conjugate of $\text{Hess}_{\mathfrak{q},x}$ by the isomorphism Π_x^{agC} , with inverse Π_x^{elC} . Hence, one Hessian is surjective if and only if the other one is. \square

In the blow-up, we can also rephrase non-degeneracy of stationary points in terms of the Hessian. Rather than stating the exact analogue of Lemma 5.30, let us emphasize the following result:

Lemma 5.31. *Let x be a stationary point of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$. Then, x is non-degenerate $\iff \text{Hess}_{\mathfrak{q},x}^{\tilde{g},\sigma}$ is injective $\iff \text{Hess}_{\mathfrak{q},x}^{\tilde{g},\sigma}$ is bijective.*

Proof. If x is irreducible, $\text{Hess}_{\mathfrak{q},x}^{\tilde{g},\sigma}$ is conjugate to $\text{Hess}_{\mathfrak{q},x}^{\tilde{g}}$ via the blow-down (which identifies $\mathcal{K}_{\text{agC},\sigma}$ with \mathcal{K}^{agC}) and we may apply Lemmas 5.19 and 5.30.

If x is a reducible stationary point in the blow-up, we have $\text{Hess}_{\mathfrak{q},x}^{\tilde{g},\sigma} = \text{Hess}_{\mathfrak{q},x}^{\sigma}$; see Remark 5.21. The claim then follows from the analogous statement for $\text{Hess}_{\mathfrak{q},x}^{\sigma}$ in [KM07, Section 12.4]. \square

So far we have only worked with W^{σ} . The above constructions can also be phrased in terms of the quotient W^{σ}/S^1 . Given $x \in W^{\sigma}$, we will write $[x]$ for its class in W^{σ}/S^1 . Note that the bundles $\mathcal{K}_j^{\text{agC},\sigma}$ are S^1 -invariant and therefore the space $\mathcal{K}_{j,x}^{\text{agC},\sigma}$ is canonically identified with the tangent space at $[x]$ in the L_j^2 completion of the tangent bundle to W^{σ}/S^1 . The vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ on W_k^{σ} is S^1 -invariant and takes values in $\mathcal{K}_{k-1}^{\text{agC},\sigma}$. Thus, it descends to a vector field, denoted $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$, on W_k^{σ}/S^1 .

Lemma 5.32. *We have the following identifications, given by composing global Coulomb projection with projection to the quotient by S^1 :*

$$(118) \quad \{\text{stationary points of } \mathcal{X}_{\mathfrak{q}}^{\sigma}\} / \mathcal{G}_{k+1} \xrightarrow{\cong} \{\text{stationary points of } \mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}\}$$

and

$$(119) \quad \{\text{trajectories of } \mathcal{X}_{\mathfrak{q}}^{\sigma}\} / \mathcal{G}_{k+1} \xrightarrow{\cong} \{\text{trajectories of } \mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}\}.$$

Proof. This is immediate from (88) and (89). \square

Let x be a stationary point of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$, so that $[x]$ is a stationary point of $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$. Since $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$ is a section of $\mathcal{K}_{k-1}^{\text{agC},\sigma}$, under the identification of $\mathcal{K}_{k-1,x}^{\text{agC},\sigma}$ with the L_{k-1}^2 -completion of $T_{[x]}(W_k^{\sigma}/S^1)$, the derivative $\mathcal{D}_{[x]}^{\sigma} \mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma} := \mathcal{D}_x^{\sigma} \mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma} = \mathcal{D}_x^{\tilde{g},\sigma} \mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ takes values in $\mathcal{K}_{k-1,x}^{\text{agC},\sigma}$. In view of Equation (105), we have

$$\text{Hess}_{\mathfrak{q},x}^{\tilde{g},\sigma} = \mathcal{D}_{[x]}^{\sigma} \mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}.$$

The following is then a direct consequence of Lemma 5.31:

Lemma 5.33. *In terms of the identification (118), non-degeneracy of a stationary point x of $\mathcal{X}_{\mathfrak{q}}^{\sigma}$ is equivalent to the injectivity (or bijectivity) of $\mathcal{D}^{\sigma} \mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$ at the corresponding point $[\Pi^{\text{gC},\sigma}(x)] \in W^{\sigma}/S^1$.*

Just as in Lemma 5.33 we rephrased the non-degeneracy of stationary points, our next goal will be to rephrase the regularity condition on the moduli spaces of trajectories of \mathcal{X}_q^σ in terms of global Coulomb gauge; that is, re-write it as a condition on the moduli spaces of trajectories of $\mathcal{X}_q^{\text{agC},\sigma}$. This will be accomplished in Section 5.10. Before that, as preliminary steps, we will:

- Embed the moduli space of trajectories of $\mathcal{X}_q^{\text{agC},\sigma}$ into a larger space of paths in Coulomb gauge (in Section 5.7);
- Describe the tangent bundle to this path space (in Section 5.8);
- Linearize the equations that define the moduli space in Coulomb gauge (in Section 5.9).

5.7. Path spaces. Fix points $x, y \in W^\sigma$, and a smooth path γ_0 in W^σ from x to y , such that $\gamma_0(t)$ agrees with x near $-\infty$ and agrees with y near $+\infty$. Let $Z = \mathbb{R} \times Y$. By analogy with the definition of $\mathcal{C}_k^\tau(x, y)$ in (48), we define the space of four-dimensional configurations

$$\mathcal{C}_k^{\text{gC},\tau}(x, y) = \{\gamma \in \mathcal{C}_{k,\text{loc}}^{\text{gC},\tau}(\mathbb{R} \times Y) \mid \gamma - \gamma_0 \in L_k^2(Z; iT^*Z) \times L_k^2(\mathbb{R}; \mathbb{R}) \oplus L_k^2(Z; \mathbb{S}^+)\}.$$

We can write γ as a path

$$\gamma(t) = (a(t) + \alpha(t)dt, s(t), \phi(t)).$$

Recall from Section 5.2 that the condition $\gamma \in \mathcal{C}_{k,\text{loc}}^{\text{gC},\tau}(\mathbb{R} \times Y)$ means that γ is in pseudo-temporal gauge ($\alpha(t)$ is constant on each slice), and that the one-form component $a(t)$ is in the kernel of d^* , for all t . Moreover, since we are in the τ model, we must have $\|\phi(t)\|_{L^2(Y)} = 1$ and $s(t) \geq 0$ for all t .

The space $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ embeds in a Hilbert manifold $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$, defined as above but using $\tilde{\mathcal{C}}_{k,\text{loc}}^{\text{gC},\tau}(\mathbb{R} \times Y)$; that is, dropping the condition $s(t) \geq 0$.

Further, in the spirit of Section 5.2, we let $W_k^\tau(x, y)$ (respectively $\widetilde{W}_k^\tau(x, y)$) denote the subset of $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ (respectively $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$) consisting of configurations with $\alpha(t) = 0$, i.e., in temporal gauge.

The gauge group

$$\mathcal{G}_{k+1}^{\text{gC}}(Z) := \{u : \mathbb{R} \rightarrow S^1 \mid 1 - u \in L_{k+1}^2(\mathbb{R}; \mathbb{C})\}$$

acts on $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ and $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$.

Let $\mathcal{B}_k^{\text{gC},\tau}(x, y)$ denote the quotient of $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ by $\mathcal{G}_{k+1}^{\text{gC}}(Z)$. Note that $\mathcal{B}_k^{\text{gC},\tau}(x, y)$ only depends on the classes $[x]$ and $[y]$ in W^σ/S^1 , up to canonical diffeomorphism. Therefore, we will use the notation $\mathcal{B}_k^{\text{gC},\tau}([x], [y])$. The quotient of $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$ by the same gauge action is denoted $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$. One can check that $\mathcal{B}_k^{\text{gC},\tau}([x], [y])$ and $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$ are Hausdorff in the quotient topology; compare [KM07, Proposition 13.3.4].

Remark 5.34. It is important to note that given an element $\gamma \in \mathcal{C}_k^{\text{gC},\tau}(x, y)$, it cannot necessarily be moved to be in temporal gauge by an element of $\mathcal{G}_{k+1}^{\text{gC}}(Z)$. However, we can act by a four-dimensional gauge transformation in $\mathcal{G}_{k+1,\text{loc}}^{\text{gC}}(Z)$ to move γ to Coulomb gauge, but the result will land in $\mathcal{C}_k^{\text{gC},\tau}(x, uy)$, for some $u \in S^1$. Since $\mathcal{B}_k^{\text{gC},\tau}(x, y)$ is canonically identified with $\mathcal{B}_k^{\text{gC},\tau}(x, uy)$, we can still think of $[\gamma]$ as having a representative in temporal gauge.

These constructions are similar to those in Section 4.5, where we had a space $\mathcal{B}_k^\tau([x], [y])$ of configurations in $\mathcal{C}_k^\tau(x, y)$ modulo four-dimensional gauge transformations. One can also consider the corresponding Hilbert manifolds $\tilde{\mathcal{B}}_k^\tau([x], [y])$ and $\tilde{\mathcal{C}}_k^\tau(x, y)$.

We would like to relate $\tilde{\mathcal{B}}_k^\tau([x], [y])$ to $\tilde{\mathcal{B}}_k^{\text{gC}, \tau}([x], [y])$. We first define a map

$$\Pi^{\text{gC}, \tau} : \tilde{\mathcal{C}}_k^\tau(x, y) \rightarrow \tilde{\mathcal{C}}_k^{\text{gC}, \tau}(\Pi^{\text{gC}, \sigma}(x), \Pi^{\text{gC}, \sigma}(y))$$

by the formula

$$(120) \quad (a(t) + \alpha(t)dt, s(t), \phi(t)) \mapsto \Pi^{\text{gC}, \sigma}(a(t), s(t), \phi(t)) + (\mu_Y(\alpha(t))dt, 0, 0).$$

Lemma 5.35. *There is a well-defined, continuous map*

$$(121) \quad \Pi^{[\text{gC}], \tau} : \tilde{\mathcal{B}}_k^\tau([x], [y]) \rightarrow \tilde{\mathcal{B}}_k^{\text{gC}, \tau}([x], [y]), \quad [\gamma] \rightarrow [\Pi^{\text{gC}, \tau}(\gamma)].$$

This sends $\mathcal{B}_k^\tau([x], [y])$ to $\mathcal{B}_k^{\text{gC}, \tau}([x], [y])$.

Proof. Let us first check that $[\Pi^{\text{gC}, \tau}(\gamma)]$ does not depend on the choice of representative γ for the class $[\gamma]$. Indeed, suppose we change γ by a four-dimensional gauge transformation of the form $u : Z \rightarrow S^1$. If we ignore the $\alpha(t)dt$ component and write $x(t) = (a(t), s(t), \phi(t))$, we find that u acts on each $x(t)$ as the three-dimensional gauge transformation $u(t) := u|_{\{t\} \times Y}$. Write $u(t) = e^{f(t)}$ with $f(t) : Y \rightarrow i\mathbb{R}$. Using (67) and the fact that $Gd^*df = f - \mu_Y(f)$, we see that

$$\Pi^{\text{gC}, \sigma}(e^{f(t)} \cdot x(t)) = e^{\mu_Y(f(t))} \Pi^{\text{gC}, \sigma}(x(t)).$$

From here we get

$$\Pi^{[\text{gC}], \tau}(e^f \cdot \gamma) = e^{\mu_Y(f)} \cdot \Pi^{[\text{gC}], \tau}(\gamma),$$

where the gauge transformation e^f on the left is a four-dimensional gauge transformation, and $e^{\mu_Y(f)}$ is the slicewise constant gauge transformation obtained by taking the average of f in each slice. Thus, the class $[\Pi^{\text{gC}, \tau}(\gamma)]$ is unchanged.

The fact that slicewise application of $\Pi^{\text{gC}, \sigma}$ preserves the four-dimensional L_k^2 condition can be seen from the formula (67), together with the Sobolev multiplication rule on infinite cylinders [KM07, Theorem 13.2.2]. Further, averaging $\alpha(t)$ slicewise preserves the L_k^2 condition by the Cauchy-Schwarz inequality. Continuity of the map $\Pi^{[\text{gC}], \tau}$ follows from similar arguments.

Finally, since $s(t)$ (and hence the condition $s(t) \geq 0$) is preserved by global Coulomb projection, we have that $\mathcal{B}_k^\tau([x], [y])$ is mapped to $\mathcal{B}_k^{\text{gC}, \tau}([x], [y])$. \square

Observe that the map (121) is surjective but not injective. For example, suppose we have two configurations γ_1 and γ_2 that are in temporal gauge and that differ in each slice by a three-dimensional gauge transformation $u : Y \rightarrow S^1$ (non-constant in t). Then $\Pi^{[\text{gC}], \tau}([\gamma_1]) = \Pi^{[\text{gC}], \tau}([\gamma_2])$, but γ_1 and γ_2 are typically not gauge equivalent as four-dimensional configurations.

Now suppose that $[x]$ and $[y]$ are stationary points of $\mathcal{X}_q^{\text{agC}, \sigma}$. Define $M^{\text{agC}}([x], [y])$ to be the moduli space of trajectories of $\mathcal{X}_q^{\text{agC}, \sigma}$, considered as a subspace of $\mathcal{B}_k^{\text{gC}, \tau}([x], [y])$. We can also define $M^{\text{agC}, \text{red}}([x], [y])$ similarly.

Proposition 5.36. *Every trajectory of $\mathcal{X}_q^{\text{agC}, \sigma}$ on W^σ/S^1 (connecting two stationary points x and y as in Definition 4.16) is actually in $\mathcal{B}_k^{\text{gC}, \tau}([x], [y])$. Further, the map $\Pi^{[\text{gC}], \tau}$ produces a homeomorphism between the moduli space $M([x], [y])$ (consisting of gauge equivalence classes of trajectories of \mathcal{X}_q^σ) and $M^{\text{agC}}([x], [y])$.*

Proof. The equivalence between trajectories of \mathcal{X}_q^σ and $\mathcal{X}_q^{\text{agC},\sigma}$ (without the L_k^2 conditions) was already established in Lemma 5.32.

Furthermore, as noted in Section 4.5, it is proved in [KM07, Theorem 13.3.5] that every trajectory of \mathcal{X}_q^σ connecting x and y is gauge equivalent to a trajectory in $\mathcal{C}_k^\tau(x, y)$. Thus, if we have a trajectory γ of $\mathcal{X}_q^{\text{agC},\sigma}$ connecting two stationary points, we can lift it to one of \mathcal{X}_q^σ , and apply a gauge transformation to obtain a trajectory in $\mathcal{C}_k^\tau(x, y)$. Since the map $\Pi^{[\text{gC}],\tau}$ preserves the L_k^2 condition, we deduce that γ is in $\mathcal{B}_k^{\text{gC},\tau}([x], [y])$. \square

Remark 5.37. It also follows from the proof of Proposition 5.36 that $M^{\text{agC}}([x], [y])$ is identified with the space of trajectories of $\mathcal{X}_q^{\text{gC},\sigma}$ from x to y modulo the S^1 action, and equivalently, the zero set of $\mathcal{F}_q^{\text{gC},\tau}$, restricted to $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ modulo the action of $\mathcal{G}_{k+1}^{\text{gC}}(Z)$.

5.8. Four-dimensional Coulomb slices. Fix $x, y \in W^\sigma$. We aim to prove that $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$ is a Hilbert manifold, and to identify its tangent space. The discussion here will be modelled on the corresponding one for the space $\tilde{\mathcal{B}}_k^\tau([x], [y])$, following [KM07, Section 14.3].

Let us first review the analysis for $\tilde{\mathcal{B}}_k^\tau([x], [y])$. This space is the quotient of $\tilde{\mathcal{C}}_k^\tau(x, y)$ by the gauge action. The L_j^2 completion of the tangent space to $\tilde{\mathcal{C}}_k^\tau(x, y)$ at $\gamma = (a, s, \phi)$ is

$$(122) \quad \begin{aligned} \mathcal{T}_{j,\gamma}^\tau &= \{(b, r, \psi) \mid \text{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0, \forall t\} \\ &\subset L_j^2(Z; iT^*Z) \oplus L_j^2(\mathbb{R}; \mathbb{R}) \oplus L_j^2(Z; \mathbb{S}^+). \end{aligned}$$

The derivative of the gauge group action on $\tilde{\mathcal{C}}_k^\tau(x, y)$ is given by

$$\mathbf{d}_\gamma^\tau : L_{j+1}^2(Z; i\mathbb{R}) \rightarrow \mathcal{T}_{j,\gamma}^\tau, \quad \mathbf{d}_\gamma^\tau(\xi) = (-d\xi, 0, \xi\phi).$$

Kronheimer and Mrowka define a local slice for the gauge action, $\mathcal{S}_{k,\gamma}^\tau \subset \tilde{\mathcal{C}}_k^\tau(x, y)$, by the equation:

$$(123) \quad -d^*b + isr \text{Re}\langle i\phi, \psi \rangle + i|\phi|^2 \text{Re} \mu_Y(\langle i\phi, \psi \rangle) = 0.$$

By linearizing this equation, they obtain the four-dimensional local Coulomb slice which was previously mentioned in Section 4.5:

$$\mathcal{K}_{k,\gamma}^\tau = \ker \mathbf{d}_\gamma^{\tau,\dagger} \subset \mathcal{T}_{k,\gamma}^\tau,$$

where

$$(124) \quad \mathbf{d}_\gamma^{\tau,\dagger}(b, r, \psi) = -d^*b + is^2 \text{Re}\langle i\phi, \psi \rangle + i|\phi|^2 \text{Re} \mu_Y(\langle i\phi, \psi \rangle).$$

There are also L_j^2 completions $\mathcal{K}_{j,\gamma}^\tau$ for $j \leq k$. Let $\mathcal{J}_{j,\gamma}^\tau$ denote the image of \mathbf{d}_γ^τ . In [KM07, Proposition 14.3.2], Kronheimer and Mrowka prove that \mathbf{d}_γ^τ is injective with closed range, $\mathbf{d}_\gamma^{\tau,\dagger}$ is surjective, and there is a bundle decomposition

$$\mathcal{T}_j^\tau = \mathcal{J}_j^\tau \oplus \mathcal{K}_j^\tau.$$

In turn, this implies that the slices $\mathcal{S}_{k,\gamma}^\tau$ are Hilbert submanifolds of $\tilde{\mathcal{C}}_k^\tau(x, y)$. Since they provide local models for the quotient $\tilde{\mathcal{B}}_k^\tau([x], [y])$, they deduce that this quotient is a Hilbert manifold. Furthermore, its tangent space at $[\gamma]$ can be identified with $\mathcal{K}_{k,\gamma}^\tau$.

We now turn to a similar discussion in global Coulomb gauge.

The L_j^2 completion of the tangent space to $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$ at $\gamma = (a, s, \phi)$ is

$$\begin{aligned} \mathcal{T}_{j,\gamma}^{\text{gC},\tau} &= \{(b, r, \psi) \mid d\beta(t) = 0, d^*(b(t)) = 0, \operatorname{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0, \forall t\} \\ &\subset L_j^2(Z; i\mathfrak{p}^*(T^*Z)) \oplus L_j^2(\mathbb{R}; \mathbb{R}) \oplus L_j^2(Z; \mathbb{S}^+), \end{aligned}$$

where we write b as $b(t) + \beta(t)dt$.

The derivative of the action of $\mathcal{G}_{k+1}^{\text{gC}}(Z)$ on $\tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$ is

$$\mathbf{d}_\gamma^{\text{gC},\tau} : L_{j+1}^2(\mathbb{R}; i\mathbb{R}) \rightarrow \mathcal{T}_{j,\gamma}^\tau, \quad \mathbf{d}_\gamma^{\text{gC},\tau}(\xi) = \left(-\frac{d\xi}{dt}dt, 0, \xi\phi\right).$$

A suitable local slice for this action is $\mathcal{S}_{k,\gamma}^{\text{gC},\tau} \subset \tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y)$, defined by:

$$(125) \quad \frac{d\beta}{dt} + i\langle i\phi, \psi \rangle_{\tilde{g}} = 0.$$

This is already linear, so the same equation defines the corresponding linearized local slice

$$\mathcal{K}_{k,\gamma}^{\text{gC},\tau} = \ker \mathbf{d}_\gamma^{\text{gC},\tau,\dagger} \subset \mathcal{T}_{k,\gamma}^{\text{gC},\tau},$$

where

$$(126) \quad \mathbf{d}_\gamma^{\text{gC},\tau,\dagger}(b, r, \psi) = \frac{d\beta}{dt} + i\langle i\phi, \psi \rangle_{\tilde{g}}.$$

We have L_j^2 completions $\mathcal{K}_{j,\gamma}^{\text{gC},\tau}$ for $j \leq k$, and we denote the image of $\mathbf{d}_\gamma^{\text{gC},\tau}$ by $\mathcal{J}_{j,\gamma}^{\text{gC},\tau}$.

Lemma 5.38. *The operator $\mathbf{d}_\gamma^{\text{gC},\tau}$ is injective with closed range, $\mathbf{d}_\gamma^{\text{gC},\tau,\dagger}$ is surjective, and there is a bundle decomposition*

$$\mathcal{T}_j^{\text{gC},\tau} = \mathcal{J}_j^{\text{gC},\tau} \oplus \mathcal{K}_j^{\text{gC},\tau}.$$

Further,

$$\mathcal{J}_j^{\text{gC},\tau} = \mathcal{J}_0^{\text{gC},\tau} \cap \mathcal{T}_j^\tau.$$

The proof of Lemma 5.38 is similar to [KM07, Proposition 14.3.2], so we omit it.

We deduce that $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$ is a Hilbert manifold, locally modelled on the slices $\mathcal{S}_{k,\gamma}^{\text{gC},\tau}$.

The tangent space to $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$ at γ can be identified with $\mathcal{K}_{k,\gamma}^{\text{gC},\tau}$.

Let us now take a quick look at the map $\Pi^{[\text{gC}],\tau}$ from (121). With respect to the smooth structures we just defined, this map is continuously differentiable. Its derivative

$$(127) \quad (\Pi_*^{[\text{gC}],\tau})_{[\gamma]} : \mathcal{K}_{j,[\gamma]}^\tau \longrightarrow \mathcal{K}_{j,\Pi^{[\text{gC}],\tau}([\gamma])}^{\text{gC},\tau}$$

is given by the formula

$$(b(t) + \beta(t)dt, r(t), \psi(t)) \mapsto \Pi_{\mathcal{K}_j^{\text{gC},\tau}} \left((\Pi_*^{\text{gC},\sigma})_{(a(t), s(t), \phi(t))} (b(t), r(t), \psi(t)) + (\mu_Y(\beta(t))dt, 0, 0) \right),$$

where $\Pi_{\mathcal{K}_j^{\text{gC},\tau}}$ denotes the projection onto $\mathcal{K}_j^{\text{gC},\tau}$ with kernel $\mathcal{J}_j^{\text{gC},\tau}$.

5.9. The linearized equations. Let x, y be non-degenerate stationary points of \mathcal{X}_q^σ . Recall from Section 4 that the moduli space $M(x, y)$ of perturbed Seiberg-Witten trajectories can be described as the zero set of the section

$$\mathcal{F}_q^\tau : \mathcal{C}_k^\tau(x, y) \rightarrow \mathcal{V}_{k-1}^\tau(Z),$$

modulo gauge. In temporal gauge, we have $\mathcal{F}_q^\tau = \frac{d}{dt} + \mathcal{X}_q^\sigma$.

Recall that $\mathcal{V}_{k-1}^\tau(Z)$ is a bundle over the Hilbert manifold $\tilde{\mathcal{C}}_k^\tau(x, y)$, and we have extended the section \mathcal{F}_q^τ . In order to understand the local structure of $M(x, y)$, one needs to study the derivative $\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau$ at paths $\gamma \in M(x, y)$. Further, to be able to define gradings and orientations later, one needs to understand $\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau$ when γ is not a trajectory in $M(x, y)$. Much like using the notation \mathcal{D}^σ to clarify derivatives in the blow-up, we write \mathcal{D}^τ to mean that the derivatives are taken with respect to four-dimensional configurations (as opposed to three-dimensional derivatives slicewise). The relevant properties of $\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau$ are analyzed by Kronheimer and Mrowka in [KM07, Section 14.4]. We will sketch their results, and then do a similar analysis in global Coulomb gauge, with an eye towards the local structure of the moduli spaces $M^{\text{agC}}([x], [y])$.

Fix $\gamma \in \tilde{\mathcal{C}}_k^\tau(x, y)$, and assume that γ is in temporal gauge. Recall the definition of the tangent space $\mathcal{T}_{j,\gamma}^\tau$ from (122). Following [KM07, Section 14.4], we write an element of $\mathcal{T}_{j,\gamma}^\tau$ as (V, β) , where $V(t) = (b(t), r(t), \psi(t))$ is a path in (a completion of) $\mathcal{T}^\sigma(Y)$, and $\beta = (\beta(t))$ is the path in $L^2(Y; i\mathbb{R})$ that gives the dt component of the connection. Vectors in $\mathcal{T}^\sigma(Y)$ can be differentiated along paths using the covariant derivative

$$\frac{D^\sigma}{dt} V = \left(\frac{db}{dt}, \frac{dr}{dt}, \Pi_{\phi(t)}^\perp \frac{d\psi}{dt} \right),$$

where $\Pi_{\phi(t)}^\perp$ denotes the L^2 projection to the orthogonal complement of $\phi(t)$. We can then write the derivative $\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau$ explicitly:

$$\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau(V, \beta) = \frac{D^\sigma}{dt} V + \mathcal{D}^\sigma \mathcal{X}_q^\sigma(V) + \mathbf{d}_{\gamma(t)}^\sigma \beta,$$

where \mathbf{d}^σ is as in (101).

Let us also recall from Section 5.8 that the L_j^2 completion of the tangent space to $\tilde{\mathcal{B}}_k^\tau([x], [y])$ at γ is the space $\mathcal{K}_{j,\gamma}^\tau = \ker \mathbf{d}_\gamma^{\tau,\dagger}$. With respect to the decomposition (V, β) , the map $\mathbf{d}_\gamma^{\tau,\dagger} : \mathcal{T}_{j,\gamma}^\tau \rightarrow L_{j-1}^2(Z; i\mathbb{R})$ is given by

$$\mathbf{d}_\gamma^{\tau,\dagger}(V, \beta) = \frac{d\beta}{dt} + \mathbf{d}_{\gamma(t)}^{\sigma,\dagger}(V),$$

with $\mathbf{d}^{\sigma,\dagger}$ as in (102).

The local structure of the moduli space $M([x], [y])$ is governed by the operator:

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau} : \mathcal{K}_{j,\gamma}^\tau \rightarrow \mathcal{V}_{j-1,\gamma}^\tau(Z).$$

Kronheimer and Mrowka establish the following:

Proposition 5.39 (Proposition 14.4.3 in [KM07]). *For $1 \leq j \leq k$, the operator $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau}$ is Fredholm and the index is independent of j .*

Sketch of proof. The main tool is Proposition 14.2.1 in [KM07], which gives a Fredholmness criterion for differential operators on infinite cylinders. Specifically, it deals with operators

of the form

$$(128) \quad Q = \frac{d}{dt} + L_0 + h_t : L_1^2(Z; p^*E) \rightarrow L^2(Z; p^*E),$$

where $E \rightarrow Y$ is a vector bundle, L_0 is a first order, self-adjoint elliptic operator acting on sections of E , and h_t is a time-dependent bounded operator on $L^2(Y; E)$, varying continuously in the operator norm topology, and assumed to be constant h_\pm near the ends of the cylinder Z . The proposition says that if $L_0 + h_\pm$ are hyperbolic (i.e., their spectrum is disjoint from the imaginary axis), then Q is Fredholm, with index given by the spectral flow of the family $\{L_0 + h_t\}$.

Proposition 14.2.1 in [KM07] cannot be applied directly to the operator $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau}$, because its domain is the local Coulomb slice, rather than the space of all sections of a vector bundle. The remedy is to enlarge the operator, much as in the proof of Lemma 5.19. We take the direct sum

$$(129) \quad Q_\gamma = \mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau \oplus \mathbf{d}_\gamma^{\tau,\dagger} : \mathcal{T}_{j,\gamma}^\tau \rightarrow \mathcal{V}_{j-1,\gamma}^\tau \oplus L_{j-1}^2(Z; i\mathbb{R}),$$

previously considered in (49). We can write

$$(130) \quad Q_\gamma = \frac{D^\sigma}{dt} + \begin{pmatrix} \mathcal{D}_{\gamma(t)}^\sigma \mathcal{X}_q^\sigma & \mathbf{d}_{\gamma(t)}^\sigma \\ \mathbf{d}_{\gamma(t)}^{\sigma,\dagger} & 0 \end{pmatrix}.$$

Further, to deal with the condition $\operatorname{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0$ that defines $\mathcal{T}_{j,\gamma}^\tau$, we proceed as in Section 5.5.2: We combine the r and ψ components into

$$\psi(t) = \psi(t) + r(t)\phi(t).$$

Now Q_γ is of the desired form (128), with E being the bundle $iT^*Y \oplus \mathbb{S} \oplus i\mathbb{R}$ and L_0 having the block form

$$(131) \quad L_0 = \begin{pmatrix} *d & 0 & -d \\ 0 & D & 0 \\ -d^* & 0 & 0 \end{pmatrix}.$$

In fact, $L_0 + h_t$ is a slight variant of the extended Hessian from (100). On the two ends of the cylinder Z we see exactly the extended Hessians at the non-degenerate stationary points x and y . Recall that at a stationary point (say $x = (a, s, \phi)$), the extended Hessian is the operator

$$\widehat{\operatorname{Hess}}_{q,x}^\sigma = \begin{pmatrix} 0 & 0 & \mathbf{d}_x^\sigma \\ 0 & \operatorname{Hess}_{q,x}^\sigma & 0 \\ \mathbf{d}_x^{\sigma,\dagger} & 0 & 0 \end{pmatrix}.$$

This is written in block form with respect to the decomposition $L_j^2(Y; E) = \mathcal{J}_j^\sigma \oplus \mathcal{K}_j^\sigma \oplus L_j^2(Y; i\mathbb{R})$; compare Sections 5.5.1 and 5.5.2.

The extended Hessians at x and y are invertible by the non-degeneracy assumption; moreover, if we had not used the blow-up, they would be self-adjoint and therefore have real spectrum. Given the use of ψ imposed from the blow-up construction, the limiting operators are no longer self-adjoint. Nevertheless, they have real spectrum by [KM07, Lemma 12.4.3]. It follows that they are hyperbolic, and the hypotheses of [KM07, Proposition 14.2.1] apply. We obtain that Q_γ is Fredholm as an operator from L_1^2 to L^2 . Further, if $L_0 + h_\pm$ are k -ASAFOE and h_t is compact as a map from L_j^2 to L_{j-1}^2 for paths constant near the endpoints, then using elliptic estimates, the result is extended to show that Q_γ as a map from L_j^2 to L_{j-1}^2 is also Fredholm with the same index. As discussed, strictly speaking, we chose the

path γ to be constant near the ends. That case implies the general case by the argument at the end of [KM07, proof of Theorem 14.4.2]. Finally, Fredholmness of Q_γ implies that $(\mathcal{D}_{[\gamma]}^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau}$ is also Fredholm, of the same index; this follows from the decomposition (129), together with the surjectivity of $\mathbf{d}_\gamma^{\tau,\dagger}$ (cf. Proposition 14.3.2 in [KM07]). \square

We now move to Coulomb gauge. Consider the moduli space $M^{\text{agC}}([x], [y]) \subset \tilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y])$, as defined before Proposition 5.36. Recall from Remark 5.37 that we can describe $M^{\text{agC}}([x], [y])$ as the zero set of

$$\mathcal{F}_q^{\text{gC},\tau} : \tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y) \rightarrow \mathcal{V}_{k-1}^{\text{gC},\tau}(Z),$$

restricted to $\mathcal{C}_k^{\text{gC},\tau}(x, y)$ and modulo the action of $\mathcal{G}_{k+1}^{\text{gC}}(Z)$.

Fix $\gamma \in \mathcal{C}_k^{\text{gC},\tau}(x, y)$ in temporal gauge (that is, in $W_k^\tau(x, y)$, cf. Section 5.7). To understand the local structure of $M^{\text{agC}}([x], [y])$, we need to study the derivative of the section $\mathcal{F}_q^{\text{gC},\tau}$. There are different ways of doing this, depending on what covariant derivative we choose. The most natural choice is to use a covariant derivative that involves the \tilde{g} -metric. Specifically, for $1 \leq j \leq k$, we write

$$(132) \quad \mathcal{D}_\gamma^{\tilde{g},\tau} \mathcal{F}_q^{\text{gC},\tau} : \mathcal{T}_{j,\gamma}^{\text{gC},\tau} \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\tau},$$

$$\mathcal{D}_\gamma^{\tilde{g},\tau} \mathcal{F}_q^{\text{gC},\tau}(V, \beta) = \frac{D^{\tilde{g},\sigma}}{dt} V + (\mathcal{D}_{\gamma(t)}^{\tilde{g},\sigma} \mathcal{X}_q^{\text{gC},\sigma})(V) + \mathbf{d}_{\gamma(t)}^{\text{gC},\sigma} \beta,$$

where

$$(133) \quad \mathbf{d}_x^{\text{gC},\sigma} : i\mathbb{R} \rightarrow \mathcal{T}_{j,x}^{\text{gC},\sigma}, \quad \mathbf{d}_x^{\text{gC},\sigma}(\xi) = (0, 0, \xi\phi)$$

is the restriction of the operator $\mathbf{d}_x^\sigma : L_j^2(Y; i\mathbb{R}) \rightarrow \mathcal{T}_{j,x}^\sigma$ from (101) to constant functions. Also, $\frac{D^{\tilde{g},\sigma}}{dt}$ represents the covariant derivative (along a path) for the connection $\mathcal{D}^{\tilde{g},\sigma}$ from (104), that is,

$$(134) \quad \frac{D^{\tilde{g},\sigma}}{dt} V = \Pi_*^{\text{gC},\sigma} \circ \mathcal{D}_{\Pi^{\text{elC},\sigma}(\frac{d\gamma}{dt})}^\sigma (\Pi^{\text{elC},\sigma}(V)).$$

We could also use the covariant derivative coming from the L^2 metric. Define

$$\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau} : \mathcal{T}_{j,\gamma}^{\text{gC},\tau} \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\tau},$$

by

$$(135) \quad \mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau}(V, \beta) = \frac{D^\sigma}{dt} V + (\mathcal{D}_{\gamma(t)}^\sigma \mathcal{X}_q^{\text{gC},\sigma})(V) + \mathbf{d}_{\gamma(t)}^{\text{gC},\sigma} \beta.$$

While the operator $\mathcal{D}_\gamma^{\tilde{g},\tau} \mathcal{F}_q^{\text{gC},\tau}$ seems more natural, it will be easier to work with $\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau}$ later on, so we will focus on the latter. When γ is a flow trajectory, then the two operators coincide.

As discussed in Section 5.8, the tangent space to $\tilde{\mathcal{B}}^{\text{gC},\tau}([x], [y])$ has completions $\mathcal{K}_{j,\gamma}^{\text{gC},\tau} = \ker \mathbf{d}_\gamma^{\text{gC},\tau,\dagger}$. From (126) we can write

$$\mathbf{d}_\gamma^{\text{gC},\tau,\dagger} = \frac{d\beta}{dt} + \mathbf{d}_{\gamma(t)}^{\text{gC},\sigma,\dagger}(V),$$

where, at $x = (a, s, \phi) \in W^\sigma$,

$$(136) \quad \mathbf{d}_x^{\text{gC},\sigma,\dagger} : \mathcal{T}_{j,x}^{\text{gC},\sigma} \rightarrow i\mathbb{R}, \quad \mathbf{d}_x^{\text{gC},\sigma,\dagger}(b, r, \psi) = i\langle i\phi, \psi \rangle_{\tilde{g}}.$$

Note that $\mathbf{d}^{\text{gC},\sigma,\tilde{\dagger}}$ appears as part of the operator $\mathbf{d}^{\sigma,\tilde{\dagger}} : \mathcal{T}_{j,x}^\sigma \rightarrow L_{j-1}^2(Y; i\mathbb{R})$ from (110). Precisely, we decompose the domain of $\mathbf{d}^{\sigma,\tilde{\dagger}}$ as follows:

$$\mathcal{T}_{j,x}^\sigma = \mathcal{K}_{j,x}^{\text{agC},\sigma} \oplus \mathcal{J}_{j,x}^{\text{gC},\sigma} \oplus \mathcal{J}_{j,x}^{\circ,\sigma} = \mathcal{T}_{j,x}^{\text{gC},\sigma} \oplus \mathcal{J}_{j,x}^{\circ,\sigma}.$$

Then, with respect to the last decomposition (where we combine the first two summands into one), we have

$$(137) \quad \mathbf{d}^{\sigma,\tilde{\dagger}} = \begin{pmatrix} 0 & -d^* \\ \mathbf{d}^{\text{gC},\sigma,\tilde{\dagger}} & 0 \end{pmatrix} : \mathcal{T}_{j,x}^{\text{gC},\sigma} \oplus \mathcal{J}_{j,x}^{\circ,\sigma} \rightarrow (\text{im } d^*)_{j-1} \oplus i\mathbb{R}.$$

Returning to our operator $\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau}$, we consider its restriction to $\mathcal{K}_{j,\gamma}^{\text{gC},\tau}$. The surjectivity of this restriction would imply that $M^{\text{agC}}([x], [y])$ is a smooth manifold near $[\gamma]$, and the dimension of its kernel would give the dimension of $M^{\text{agC}}([x], [y])$. As a preliminary step towards this discussion, we establish the following:

Proposition 5.40. *Let $x, y \in W_k^\sigma$ be non-degenerate stationary points of \mathcal{X}_q^σ (and hence of $\mathcal{X}_q^{\text{agC},\sigma}$). Pick a path $\gamma \in W_k^\tau(x, y)$. Then, for $j \leq k$, the operator*

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC},\tau}} : \mathcal{K}_{j,\gamma}^{\text{gC},\tau} \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\tau}$$

is Fredholm. Moreover, the Fredholm index is the same as that of the operator

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau} : \mathcal{K}_{j,\gamma}^\tau \rightarrow \mathcal{V}_{j-1,\gamma}^\tau.$$

Proof. To establish the Fredholm property, the arguments are similar to those in the proof of Proposition 5.39, with some modifications.

We extend our operator $\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau}$ so that it acts on sections of a vector bundle. First, define

$$(138) \quad Q_\gamma^{\text{gC}} = \mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau} \oplus \mathbf{d}_\gamma^{\text{gC},\tau,\tilde{\dagger}} : \mathcal{T}_{j,\gamma}^{\text{gC},\tau} \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\tau} \oplus L_{j-1}^2(\mathbb{R}; i\mathbb{R}).$$

The codomain of $Q_\gamma^{\text{gC},\tilde{g},\sigma}$ can be identified with $\mathcal{T}_{j-1,\gamma}^{\text{gC},\tau}$. However, even after writing $\psi(t) = \psi(t) + r(t)\phi(t)$ as before, $\mathcal{T}_{j,\gamma}^{\text{gC},\tau}$ is still not the space of all sections of a vector bundle. Indeed, for $(b(t) + \beta(t)dt, \psi(t)) \in \mathcal{T}_{j,\gamma}^{\text{gC},\tau}$, we still have the conditions $d(\beta(t)) = 0$ and $d^*(b(t)) = 0$.

Thus, we need to extend the operator once more. Consider the linear operator

$$(139) \quad R = \frac{d}{dt} + \begin{pmatrix} 0 & -d \\ -d^* & 0 \end{pmatrix} : (\text{im } d_{(0)} \oplus \text{im } d_{(1)}^*) \rightarrow (\text{im } d_{(0)} \oplus \text{im } d_{(1)}^*)$$

where the subscript (p) with $p \in \{0, 1\}$ denotes the imaginary p -forms on which the respective operator acts, on each slice $\{t\} \times Y$. Precisely, $\text{im } d_{(1)}^*$ is the subset of $L_j^2(Z; i\mathbb{R})$ consisting of functions that integrate to zero slicewise, and $\text{im } d_{(0)} = \ker d_{(1)}$, since Y is a rational homology sphere.

Next, decompose $\mathcal{T}_{j,\gamma}^\tau$ as

$$(140) \quad \mathcal{T}_{j,\gamma}^\tau = \mathcal{T}_{j,\gamma}^{\text{gC},\tau} \oplus (\mathcal{J}_{j,\gamma}^{\circ,\tau} \oplus \text{im } d_{(1)}^*),$$

where $\mathcal{J}_{j,\gamma}^{\circ,\tau}$ consists of time-dependent elements of the spaces $\mathcal{J}_{j,\gamma(t)}^{\circ,\sigma}$ from (76). Note that there is a natural identification

$$\Psi : \text{im } d_{(0)} \rightarrow \mathcal{J}_{j,\gamma}^{\circ,\tau}$$

given at each time t by the formula

$$\Psi(-d\xi, 0, 0) = (-d\xi, 0, \xi \cdot \phi(t)),$$

where $\phi(t)$ is the spinor component of $\gamma(t)$. If we conjugate the operator R by Ψ on the $\text{im } d_{(0)}$ summand, we obtain an operator

$$(141) \quad \widehat{R} = \Psi \circ \frac{d}{dt} \circ \Psi^{-1} + \begin{pmatrix} 0 & \mathbf{d}^\sigma \\ -d^* & 0 \end{pmatrix} : (\mathcal{J}_{j,\gamma}^{\circ,\tau} \oplus \text{im } d_{(1)}^*) \rightarrow (\mathcal{J}_{j,\gamma}^{\circ,\tau} \oplus \text{im } d_{(1)}^*).$$

Observe that $\Psi \circ \frac{d}{dt} \circ \Psi^{-1}(-d\xi, 0, \xi\phi) = (-\frac{d\xi}{dt}, 0, \frac{d\xi}{dt}\phi)$.

Define the new extension of Q_γ^{gC} to be

$$(142) \quad \widehat{Q}_\gamma^{\text{gC}} = \begin{pmatrix} Q_\gamma^{\text{gC}} & 0 \\ 0 & \widehat{R} \end{pmatrix} : \mathcal{T}_{j,\gamma}^\tau \rightarrow \mathcal{T}_{j-1,\gamma}^\tau,$$

with respect to the decomposition (140).

Instead of (140) we could consider the decomposition $\mathcal{T}_{j,\gamma}^\tau = \mathcal{V}_{j,\gamma}^\tau \oplus L_j^2(Z; i\mathbb{R})$, where we recall that \mathcal{V}^τ is the space of four-dimensional configurations with trivial dt component that are slice-wise in \mathcal{T}^σ . With respect to this new decomposition, we can write

$$(143) \quad \widehat{Q}_\gamma^{\text{gC}} = \frac{D^\sigma}{dt} + \begin{pmatrix} M & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} H & \mathbf{d}_{\gamma(t)}^\sigma \\ \mathbf{d}_{\gamma(t)}^{\sigma,\dagger} & 0 \end{pmatrix}.$$

Here, for a fixed time t , the operator M acts by zero on $\mathcal{T}_{j,\gamma(t)}^{\text{gC},\sigma} \subset \mathcal{T}_{j,\gamma(t)}^\sigma$ and equals the difference $\Psi \circ \frac{d}{dt} \circ \Psi^{-1} - \frac{D^\sigma}{dt}$ when applied to elements of $\mathcal{J}_{j,\gamma(t)}^{\circ,\sigma} \subset \mathcal{T}_{j,\gamma(t)}^\sigma$; that is,

$$M(-d\xi, 0, \xi\phi) = \Pi_\phi^\perp(0, 0, \xi\frac{d\phi}{dt}).$$

Furthermore, with respect to the decomposition $\mathcal{T}_{j,\gamma(t)}^\sigma = \mathcal{T}_{j,\gamma(t)}^{\text{gC},\sigma} \oplus \mathcal{J}_{j,\gamma(t)}^{\circ,\sigma}$, the operator H from (143) is given by

$$(144) \quad H = \begin{pmatrix} \mathcal{D}_{\gamma(t)}^\sigma \mathcal{X}_q^{\text{gC},\sigma} & 0 \\ 0 & 0 \end{pmatrix}.$$

Note that the third term in (143) is exactly $\mathcal{H}_{\gamma(t)}^\sigma$ from Lemma 5.26.

Using Lemma 5.26 and the arguments in the proof of Proposition 5.39, we obtain that $\widehat{Q}_\gamma^{\text{gC}}$ is of the form (128), with the bundle $E = iT^*Y \oplus \mathbb{S} \oplus i\mathbb{R}$ just as for Q_γ . Furthermore, the differential part L_0 in $\widehat{Q}_\gamma^{\text{gC}}$ is the same L_0 that appeared in (131) for Q_γ . We write

$$(145) \quad \widehat{Q}_\gamma^{\text{gC}} = \frac{d}{dt} + L_0 + \hat{h}_t^{\text{gC}}.$$

It's important to note that $L_0 + \hat{h}_t^{\text{gC}}$ is not exactly the operator \mathcal{H}_x^σ arising from Lemma 5.26, as \hat{h}_t^{gC} also has terms coming from the time derivative of ϕ . In the limit, as $t \rightarrow \pm\infty$ (i.e., as we approach the stationary points x and y), we do have from (144) and Lemma 5.26 that

$$(146) \quad L_0 + \hat{h}_{\pm\infty}^{\text{gC}} = \widehat{\text{Hess}}_{q,\pm}^{\tilde{g},\sigma},$$

where $\widehat{\text{Hess}}_q^{\tilde{g},\sigma}$ is the extended Hessian from (109). The extended Hessian is hyperbolic at non-degenerate stationary points; indeed, it has real spectrum by Lemma 5.27. Using the properties of \mathcal{H}_x^σ established in Lemma 5.26, we can apply Proposition 14.2.1 in [KM07] together with the arguments in [KM07, p.256] (mentioned in the proof of Proposition 5.39) and deduce that $\widehat{Q}_\gamma^{\text{gC}}$ is Fredholm. As before, first we establish Fredholmness as a map from L_1^2 to L^2 for paths γ which are constant near the endpoints, and then we extend this to

operators from L_j^2 to L_{j-1}^2 for $1 \leq j \leq k$ and for all paths. Again the index is independent of j .

We claim that $\widehat{Q}_\gamma^{\text{gC}}$ has the same Fredholm index as the operator Q_γ from (129). Indeed, we can relate them by a continuous family of Fredholm operators, in two steps, along the lines of Section 5.5.5. First, we interpolate linearly from Q_γ to the operator

$$Q_\gamma^{\text{sp}} := \frac{D^\sigma}{dt} + \begin{pmatrix} \mathcal{D}_{\gamma(t)}^\sigma \mathcal{X}_q^\sigma & \mathbf{d}_{\gamma(t)}^\sigma \\ \mathbf{d}_{\gamma(t)}^{\text{sp}, \sigma, \dagger} & 0 \end{pmatrix},$$

which differs from Q_γ by having $\mathbf{d}_{\gamma(t)}^{\text{sp}, \sigma, \dagger}$ instead of $\mathbf{d}_{\gamma(t)}^{\sigma, \dagger}$. Second, we use the family of metrics g_ρ from Section 5.5.5 to define operators \widehat{Q}_γ^ρ similar to $\widehat{Q}_\gamma^{\text{gC}}$, such that at $\rho = 0$ we have

$$\widehat{Q}_\gamma^0 = Q_\gamma^{\text{sp}}$$

and at $\rho = 1$ we have

$$\widehat{Q}_\gamma^1 = \widehat{Q}_\gamma^{\text{gC}}.$$

To see that the operators considered during these two interpolations are Fredholm and have the same index, observe that they are all of the form (128), with the same differential part L_0 . Further, at the endpoints they limit to the interpolations between extended Hessians that appeared in Lemmas 5.28 and 5.29, and which were shown there to be invertible with real spectrum.

We have now shown that Q_γ and $\widehat{Q}_\gamma^{\text{gC}}$ have the same index. To go back from $\widehat{Q}_\gamma^{\text{gC}}$ to Q_γ^{gC} , note that the operator \widehat{R} defined in (141) is bijective. Indeed, it is conjugate to the operator R from (139), so it suffices to check that R is bijective. We have $R = \frac{d}{dt} + A$, where

$$A = \begin{pmatrix} 0 & -d \\ -d^* & 0 \end{pmatrix}$$

is a self-adjoint operator acting slicewise. Observe that A is invertible because $b_1(Y) = 0$, and we can find a complete orthonormal system of eigenvectors $\{a_n\}$, with $A(a_n) = \lambda_n a_n$, $\lambda_n \neq 0$. An element in the kernel of R would be of the form

$$\sum_n c_n e^{-\lambda_n t} a_n, \quad c_n \in \mathbb{R}.$$

However, any such nonzero element increases exponentially at one of the ends of the cylinder Z , and hence is not in L_j^2 . Thus, $\ker(R) = 0$, and by self-adjointness, $\text{coker}(R) = 0$.

Since \widehat{R} is bijective, from (142) we deduce that Q_γ^{gC} is Fredholm, of the same index as $\widehat{Q}_\gamma^{\text{gC}}$ (and hence as Q_γ). In the proof of Proposition 5.39 we saw that Q_γ has the same index as $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{j,\gamma}^\tau}$. In a similar fashion, since $\mathbf{d}_{\gamma}^{\text{gC}, \tau, \dagger}$ is surjective by Lemma 5.38, we see that $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC}, \tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC}, \tau}}$ is Fredholm and of the same index as $Q_\gamma^{\text{gC}, \tilde{g}, \sigma}$ by the same arguments as in [KM07, Proposition 14.4.3]. The conclusion follows. \square

We summarize the results obtained from the proof of Proposition 5.40.

Lemma 5.41. *Under the hypotheses of Proposition 5.40,*

(a) *the operators*

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC}, \tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC}, \tau}}, \quad Q_\gamma^{\text{gC}}, \quad \widehat{Q}_\gamma^{\text{gC}}$$

are all Fredholm of the same index;

(b) one of the operators

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC},\tau}}, Q_\gamma^{\text{gC}}, \widehat{Q}_\gamma^{\text{gC}}$$

is surjective if and only if any of the two others is surjective.

Proof. To establish the relation between $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC},\tau}}$ and $Q_\gamma^{\text{gC},\tilde{g},\sigma}$ we use Lemma 5.38, which gives the surjectivity of $\mathbf{d}_\gamma^{\text{gC},\tau,\dagger}$. To establish the relation between Q_γ^{gC} and $\widehat{Q}_\gamma^{\text{gC}}$, we use the block form (142) and the bijectivity of \widehat{R} . \square

5.10. Non-degeneracy of trajectories in Coulomb gauge. Let us now suppose that γ is a trajectory of \mathcal{X}_q^σ on $\mathcal{C}_k^\sigma(Y)$ between two non-degenerate stationary points $x, y \in \mathcal{C}^\sigma(Y)$. Recall from Definition 4.18 that the moduli space $M(x, y)$ is regular at γ if Q_γ is surjective. By [KM07, Propositions 14.3.2 and 14.4.3], this is equivalent to the operator

$$(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau} : \mathcal{K}_{k,\gamma}^\tau \rightarrow \mathcal{V}_{k-1,\gamma}^\tau$$

being surjective. (Compare the proof of Proposition 5.39.) We rephrase this condition in Coulomb gauge:

Proposition 5.42. *Consider a path $\gamma \in \mathcal{C}_k^\tau(x, y)$ in temporal gauge. Write $x^b = \Pi^{\text{gC},\sigma}(x)$, $y^b = \Pi^{\text{gC},\sigma}(y)$ and $\gamma^b = \Pi^{\text{gC},\tau}(\gamma)$. Then:*

(a) *The operators $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau} : \mathcal{K}_{k,\gamma}^\tau \rightarrow \mathcal{V}_{k-1,\gamma}^\tau$ and $(\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{k,\gamma^b}^{\text{gC},\tau}} : \mathcal{K}_{k,\gamma^b}^{\text{gC},\tau} \rightarrow \mathcal{V}_{k-1,\gamma^b}^{\text{gC},\tau}$*

have the same Fredholm index.

(b) *Suppose that γ is a trajectory of \mathcal{X}_q^σ , so that $[\gamma^b] \in \mathcal{B}_k^{\text{gC},\tau}([x^b], [y^b])$ is a trajectory of $\mathcal{X}_q^{\text{agC},\sigma}$. If $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau}$ is surjective, then so is $(\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{k,\gamma^b}^{\text{gC},\tau}}$.*

Proof. (a) We can interpolate between the paths γ and γ^b by a continuous family of paths $\{\gamma_s\}_{s \in [0,1]}$, with the endpoints of γ_s varying on the gauge orbits of x and y . Thus, the operators Q_γ and Q_{γ^b} are part of a family of Fredholm operators Q_{γ_s} , $s \in [0,1]$. All Fredholm operators in a continuous family must have the same index. Combining this with Proposition 5.40, we obtain

$$\text{ind}((\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{k,\gamma^b}^{\text{gC},\tau}}) = \text{ind}((\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma^b}^\tau}) = \text{ind}((\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau}),$$

as desired.

(b) Recall that in Section 5.7 we defined the map

$$\Pi^{\text{gC},\tau} : \widetilde{\mathcal{C}}_k^\tau(x, y) \rightarrow \widetilde{\mathcal{C}}_k^{\text{gC},\tau}(x^b, y^b)$$

and that after dividing by the gauge groups, this induces a map between the respective quotients:

$$\Pi^{[\text{gC}],\tau} : \widetilde{\mathcal{B}}_k^\tau([x], [y]) \rightarrow \widetilde{\mathcal{B}}_k^{\text{gC},\tau}([x], [y]).$$

Note that when we work with gauge equivalence classes in $\mathcal{C}^\sigma(Y)$, we have $[x] = [x^b]$ and $[y] = [y^b]$.

The derivative of $\Pi^{[\text{gC}],\tau}$, the map

$$(\Pi_*^{[\text{gC}],\tau})_{[\gamma]} : \mathcal{K}_{k,\gamma}^\tau \longrightarrow \mathcal{K}_{k,\gamma^b}^{\text{gC},\tau},$$

was mentioned in (127).

For $j \leq k$, observe that inside the completed tangent bundle \mathcal{T}_j^τ we have the subbundle \mathcal{V}_j^τ consisting of paths $(b(t), r(t), \psi(t))$ in temporal gauge. The derivative $(\Pi_*^{\text{gC}, \tau})_\gamma$ maps \mathcal{V}_j^τ to $\mathcal{V}_j^{\text{gC}, \tau}$ by applying infinitesimal global Coulomb projection slicewise:

$$(\Pi_*^{\text{gC}, \tau})_\gamma(b, r, \psi)(t) = (\Pi_*^{\text{gC}, \sigma})_{\gamma(t)}(b(t), r(t), \psi(t)).$$

Consider the diagram

$$(147) \quad \begin{array}{ccc} \tilde{\mathcal{B}}_k^\tau([x], [y]) & \xrightarrow{\mathcal{F}_q^\tau} & \mathcal{V}_{k-1}^\tau \\ \Pi^{[\text{gC}], \tau} \downarrow & & \downarrow \Pi_*^{\text{gC}, \tau} \\ \tilde{\mathcal{B}}_k^{\text{gC}, \tau}([x^\flat], [y^\flat]) & \xrightarrow{\mathcal{F}_q^{\text{gC}, \tau}} & \mathcal{V}_{k-1}^{\text{gC}, \tau}. \end{array}$$

We claim that this diagram commutes. This can be seen as follows. Choose a representative γ of $[\gamma] \in \tilde{\mathcal{B}}_k^\tau([x], [y])$ in temporal gauge.

Since the elements of \mathcal{V}_j^τ are in temporal gauge, we can write $\mathcal{F}_q^\tau = \frac{d}{dt} + \mathcal{X}_q^\sigma$ and $\mathcal{F}_q^{\text{gC}, \tau} = \frac{d}{dt} + \mathcal{X}_q^{\text{gC}, \sigma}$.

Let us also write $\gamma^\flat(t) = \Pi_*^{\text{gC}, \sigma}(\gamma(t))$ as $g(t)\gamma(t)$, where $g(t)$ is a path of gauge transformations on Y . Since the (perturbed) CSD functional is gauge-invariant, we have that \mathcal{X}_q^σ is gauge-equivariant and thus:

$$(148) \quad \mathcal{X}_q^\sigma(\gamma^\flat(t)) = \mathcal{X}_q^\sigma(g(t)\gamma(t)) = g(t)_* \mathcal{X}_q^\sigma(\gamma(t)).$$

Moreover, direct computation shows that

$$(149) \quad (\Pi_*^{\text{gC}, \sigma})_{\gamma^\flat(t)} \circ g(t)_* = (\Pi_*^{\text{gC}, \sigma})_{\gamma(t)}.$$

Since $\mathcal{X}_q^{\text{gC}, \sigma} = \Pi_*^{\text{gC}, \sigma} \circ \mathcal{X}_q^\sigma$, applying $\Pi_*^{\text{gC}, \sigma}$ to (148) and using (149) shows that

$$(150) \quad \mathcal{X}_q^{\text{gC}, \sigma}(\gamma^\flat(t)) = (\Pi_*^{\text{gC}, \sigma})_{\gamma^\flat(t)} \mathcal{X}_q^\sigma(\gamma^\flat(t)) = (\Pi_*^{\text{gC}, \sigma})_{\gamma(t)} \mathcal{X}_q^\sigma(\gamma(t)).$$

By the chain rule, we have

$$\frac{d\gamma^\flat(t)}{dt} = \frac{d}{dt}(\Pi_*^{\text{gC}, \sigma} \circ \gamma(t)) = (\Pi_*^{\text{gC}, \sigma})_{\gamma(t)} \frac{d\gamma(t)}{dt},$$

or, for short,

$$(151) \quad \frac{d}{dt}(\gamma^\flat) = \Pi_*^{\text{gC}, \tau} \left(\frac{d}{dt} \gamma \right).$$

Using (150) and (151) we compute

$$\begin{aligned} \mathcal{F}_q^{\text{gC}, \tau} \circ \Pi^{[\text{gC}], \tau}(\gamma) &= \frac{d}{dt}(\gamma^\flat) + \mathcal{X}_q^{\text{gC}, \sigma}(\gamma^\flat) \\ &= \Pi_*^{\text{gC}, \tau} \left(\frac{d}{dt} \gamma \right) + \Pi_*^{\text{gC}, \tau}(\mathcal{X}_q^\sigma \gamma) \\ &= \Pi_*^{\text{gC}, \tau}(\mathcal{F}_q^\tau(\gamma)). \end{aligned}$$

Thus, the diagram (147) commutes.

Taking derivatives in (147) and using the fact that γ and γ^b are trajectories of the respective vector fields, we obtain the commutative diagram

$$(152) \quad \begin{array}{ccc} \mathcal{K}_{k,\gamma}^\tau & \xrightarrow{\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau} & \mathcal{V}_{k-1,\gamma}^\tau \\ \Pi_*^{\text{gC},[\tau]} \downarrow & & \downarrow \Pi_*^{\text{gC},\tau} \\ \mathcal{K}_{k,\gamma^b}^{\text{gC},\tau} & \xrightarrow{\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau}} & \mathcal{V}_{k-1,\gamma^b}^{\text{gC},\tau} \end{array}$$

Since the infinitesimal global Coulomb projection $\Pi_*^{\text{gC},\sigma}$ is surjective, we see that the right vertical arrow $\Pi_*^{\text{gC},\tau}$ in (152) is also surjective. Using the commutativity of (152), we get that if $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau}$ is surjective, then so is $(\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{k,\gamma^b}^{\text{gC},\tau}}$. \square

Recall that Proposition 5.36 guaranteed that the moduli spaces $M([x], [y])$ and $M^{\text{agC}}([x], [y])$ are homeomorphic. Proposition 5.42 says more: If $M([x], [y])$ is cut out smoothly and transversely by \mathcal{F}_q^τ , then also $M^{\text{agC}}([x], [y])$ is cut out smoothly and transversely by $\mathcal{F}_q^{\text{gC},\tau}$.

So far we have only discussed the usual moduli spaces $M([x], [y])$. Similar arguments apply to the moduli spaces $M^{\text{red}}([x], [y])$ between reducibles. The end result is that we can identify them with the corresponding moduli spaces $M^{\text{agC},\text{red}}([x], [y])$ in Coulomb gauge, and that regularity of the former implies regularity of the latter.

5.11. Gradings. From now on we will assume that we have chosen an admissible perturbation \mathfrak{q} , so that all the stationary points are non-degenerate, and the moduli spaces $M([x], [y])$ and $M^{\text{red}}([x], [y])$ are regular.

Let x, y be stationary points in W^σ . In view of Lemma 5.41, Proposition 5.42 and the definition of relative gradings in Section 4.8, we have

$$(153) \quad \begin{aligned} \text{gr}(x, y) &= \text{ind } (\mathcal{D}_\gamma^\tau \mathcal{F}_q^\tau)|_{\mathcal{K}_{k,\gamma}^\tau} = \text{ind } Q_\gamma \\ &= \text{ind } (\mathcal{D}_{\gamma^b}^\tau \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{k,\gamma^b}^{\text{gC},\tau}} = \text{ind } Q_{\gamma^b}^{\text{gC}}, \end{aligned}$$

where γ is any path from x to y (not necessarily a trajectory of \mathcal{X}_q^σ).

Thus, the relative gradings can be calculated directly in Coulomb gauge.

Recall from (58) that the relative gradings of reducible stationary points of \mathcal{X}_q with the same connection component can be computed in terms of the spectrum of the operator $D_{q,a}$; here, an eigenvector ϕ of $D_{q,a}$ has eigenvalue given by the spinorial energy of $(a, 0, \phi)$. It is thus natural to ask how this relates to the analogous setup in Coulomb gauge. The following tells us that spinorial energy and “Coulomb gauge spinorial energy” are equal.

Lemma 5.43. *Let $x = (a, s, \phi) \in W_k^\sigma$. Then,*

$$(154) \quad \text{Re} \langle \widetilde{\mathcal{X}_q^{\text{gC}}}^1(a, s, \phi), \phi \rangle_{L^2} = \text{Re} \langle \widetilde{\mathcal{X}_q}^1(a, s, \phi), \phi \rangle_{L^2}.$$

Proof. By continuity, it suffices to establish (154) for irreducibles. We have

$$\begin{aligned} \operatorname{Re}\langle \widetilde{\mathcal{X}_q^{\text{gC}}}^1(a, s, \phi), \phi \rangle_{L^2} &= \frac{1}{s} \operatorname{Re}\langle (\mathcal{X}_q^{\text{gC}})^1(a, s\phi), \phi \rangle_{L^2} \\ &= \frac{1}{s} \operatorname{Re}\langle \mathcal{X}_q^1(a, s\phi) + Gd^*(\mathcal{X}_q^0(a, s\phi))\phi, \phi \rangle_{L^2} \\ &= \frac{1}{s} \operatorname{Re}\langle \mathcal{X}_q^1(a, s\phi), \phi \rangle_{L^2} \\ &= \operatorname{Re}\langle \widetilde{\mathcal{X}_q}^1(a, s, \phi), \phi \rangle_{L^2}, \end{aligned}$$

where in the penultimate equality, we use that $\operatorname{Re}\langle Gd^*(a, s\phi)\phi, \phi \rangle_{L^2} = 0$ since $Gd^*(a, s\phi)$ is purely imaginary. \square

In light of Lemma 5.43, we do not define a separate spinorial energy in Coulomb gauge.

Consider $x = (a, 0, \phi)$ a reducible stationary point of $\mathcal{X}_q^{\text{gC}, \sigma}$. Then ϕ is an eigenvector of the operator $D_{q,a}^{\text{gC}}$, defined by

$$D_{q,a}^{\text{gC}}(\psi) := \mathcal{D}_{(a,0)}(\mathcal{X}_q^{\text{gC}})^1(0, \psi) = \widetilde{\mathcal{X}_q^{\text{gC}}}^1(a, 0, \psi).$$

Let μ be the corresponding eigenvalue. Since x is necessarily a stationary point of \mathcal{X}_q^σ as well, ϕ is an eigenvector of $D_{q,a}$ and $(a, 0)$ is a stationary point of \mathcal{X}_q . In particular, $\mathcal{D}_{(a,0)}\mathcal{X}_q(0, \phi)$ is in Coulomb gauge, and thus

$$\mu\phi = D_{q,a}(\phi) = D_{q,a}^{\text{gC}}(\phi).$$

It follows that

$$(155) \quad \Lambda_q(x) = \langle \widetilde{\mathcal{X}_q^{\text{gC}}}^1(0, \phi), \phi \rangle_{L^2} = \langle \mu\phi, \phi \rangle_{L^2} = \mu,$$

which agrees with Lemma 5.43. Note that here we do not need to take real parts, since the relevant operators are self-adjoint at stationary points.

Lemma 5.44. *Fix a reducible stationary point $(a, 0)$ of $\mathcal{X}_q^{\text{gC}}$. For each $N \in \mathbb{N}$, there exist $\omega_1, \omega_2 > 0$ such that the (finitely many) reducible stationary points of $\mathcal{X}_q^{\text{agC}, \sigma}$ which agree with $(a, 0)$ in the blow-down and have grading in the interval $[-N, N]$ are precisely the reducible stationary points with spinorial energy in the interval $[-\omega_1, \omega_2]$.*

Proof. This follows from (58) and (155). \square

5.12. The cut-down moduli spaces in Coulomb gauge. Recall that in Section 4.7, the U -map on monopole Floer homology was defined by intersecting the moduli spaces $M([x], [y])$ and $M^{\text{red}}([x], [y])$ with the zero set \mathcal{Z} of a transverse section ζ of a complex line bundle E^σ over $\mathcal{B}_k^\sigma(B_p)$. The bundle E^σ was associated to the map $\mathcal{G}_{k+1}(B_p) \rightarrow S^1$, $u \mapsto u(p)$, where $p = (t, q)$ is a point in $\mathbb{R} \times Y$. Without loss of generality, let us assume that $t = 0$.

We will need to modify this definition so that we can relate it to Coulomb gauge. First, whereas in Section 4.7 we followed [KMOS07] and used the restriction of configurations to a standard ball B_p around p , we could just as well restrict to any closed neighborhood N_p of p that is a manifold with boundary. It is convenient to take $N_p = [-1, 1] \times Y$. Consider the restriction $r : \mathcal{B}_k(N_p) \rightarrow \mathcal{B}_k(B_p)$. On the blow-up, this induces a map

$$r^\sigma : \mathcal{N} \rightarrow \mathcal{B}_k^\sigma(B_p),$$

$$r^\sigma(a, s, \phi) = (a, s \cdot \|\phi\|_{L^2(B_p)}, \phi/\|\phi\|_{L^2(B_p)}),$$

well-defined on the open subset $\mathcal{N} \subset \mathcal{B}_k^\sigma(N_p)$ consisting of configurations (a, s, ϕ) such that ϕ does not vanish identically on B_p . Note that, because of the unique continuation principle, the moduli spaces $M([x], [y])$ and $M^{\text{red}}([x], [y])$ are contained in \mathcal{N} . (Here, we identify trajectories with their restrictions to N_p , again using unique continuation.)

Let us pull back ζ under r^σ and obtain a section $(r^\sigma)^*\zeta$ of the bundle

$$(E^\sigma)' := (r^\sigma)^*E^\sigma$$

over \mathcal{N} . Intersecting the moduli spaces with the zero set of $(r^\sigma)^*\zeta$ is the same as intersecting them with \mathcal{Z} . Furthermore, we can consider any other transverse section of $(E^\sigma)'$, and let \mathcal{Z}' be its zero set. If we define the U -maps using intersections with \mathcal{Z}' , standard continuation arguments in Floer theory show that they are chain homotopic to the original ones.

Note that $(E^\sigma)'$ is associated to the map $\mathcal{G}_{k+1}(N_p) \rightarrow S^1$, $u \mapsto u(0, q)$. For our next modification, let $(E^\sigma)''$ be the complex line bundle over $\mathcal{N} \subset \mathcal{B}_k^\sigma(N_p)$ associated to the map

$$(156) \quad \mathcal{G}_{k+1}(N_p) \rightarrow S^1, \quad e^f \mapsto e^{\mu_Y(f(0, \cdot))},$$

where $f : N_p \rightarrow i\mathbb{R}$ and $\mu_Y(f(0, \cdot))$ is the average value of f on the slice $\{0\} \times Y$.

We can construct a family of bundles interpolating between $(E^\sigma)'$ and $(E^\sigma)''$, by considering the maps

$$\mathcal{G}_{k+1}(N_p) \rightarrow S^1, \quad e^f \mapsto e^{\lambda\mu_Y(f(0, \cdot)) + (1-\lambda)f(0, q)}.$$

Thus, instead of considering sections of $(E^\sigma)'$, we could define the U -maps using sections of $(E^\sigma)''$, and the results will be chain homotopic to the originals.

The third modification consists in moving from the σ model to the τ model. As explained in [KM07, Section 6.3], the two models are equivalent in the following sense. Consider the open subset $\mathcal{U} \subset \mathcal{B}_k^\sigma(N_p)$ consisting of configurations (a, s, ϕ) with $\phi|_{\{t\} \times Y} \not\equiv 0$ for all $t \in [0, 1]$. As noted in [KM07, p.463], there is a natural map

$$\varrho : \mathcal{U} \rightarrow \mathcal{B}_k^\tau(N_p).$$

By unique continuation, the moduli spaces $M([x], [y])$ and $M^{\text{red}}([x], [y])$ restricted to N_p yield configurations in \mathcal{U} , which map homeomorphically onto their image under ϱ .

There is a complex line bundle E^τ over $\mathcal{B}_k^\tau(N_p)$ associated to the same map (156) as before. When we pull it back under ϱ and then further restrict to $\mathcal{U} \cap \mathcal{N}$, we obtain the restriction of the bundle $(E^\sigma)''$ to $\mathcal{U} \cap \mathcal{N}$. Thus, we can equivalently define the U -maps using the τ model and intersecting with the zero set \mathcal{Z}^τ of a transverse section of E^τ .

We are now ready to make the connection with configurations in Coulomb gauge. This is done via the global Coulomb projection $\Pi^{[\text{gC}], \tau}$. In (121) the map $\Pi^{[\text{gC}], \tau}$ was defined for configurations on $\mathbb{R} \times Y$ with fixed asymptotics, but the same formula (120) can be applied to a configuration on $N_p = [-1, 1] \times Y$. By a slight abuse of notation, we still write $\Pi^{[\text{gC}], \tau}$ for the resulting map $\mathcal{B}_k^\tau(N_p) \rightarrow \mathcal{B}_k^{\text{gC}, \tau}(N_p)$. Here, $\mathcal{B}_k^{\text{gC}, \tau}(N_p)$ stands for the quotient of the space $\mathcal{C}_k^{\text{gC}, \tau}(N_p)$ by the gauge group $\mathcal{G}_{k+1}^{\text{gC}}(N_p)$; cf. Section 5.2.

Note that there is a natural map

$$(157) \quad \mathcal{B}_k^{\text{gC}, \tau}(N_p) \rightarrow W_{k-1/2}^\sigma/S^1, \quad [\gamma] \rightarrow [\gamma(0)],$$

where the representative γ is chosen to be in temporal gauge.

Consider the complex line bundle $E^{\text{agC},\sigma}$ over $W_{k-1/2}^\sigma/S^1$, associated to the S^1 -bundle $W_{k-1/2}^\sigma$. Pick a section ζ^{agC} transverse to the zero section, and such that the zero set \mathcal{Z}^{agC} of ζ^{agC} intersects all the moduli spaces $M^{\text{agC}}([x], [y])$ and $M^{\text{agC},\text{red}}([x], [y])$ transversely. We obtain cut-down moduli spaces in Coulomb gauge:

$$M^{\text{agC}}([x], [y]) \cap \mathcal{Z}^{\text{agC}} \quad \text{and} \quad M^{\text{agC},\text{red}}([x], [y]) \cap \mathcal{Z}^{\text{agC}}.$$

Here, we identified $M^{\text{agC}}([x], [y])$ and $M^{\text{agC},\text{red}}([x], [y])$ with their images in $W_k^\sigma/S^1 \subset W_{k-1/2}^\sigma/S^1$ at time $t = 0$, for simplicity. (See [KM07, Proposition 7.2.1] for the model unique continuation result for this case.) Alternatively, we could identify them with their images in $\mathcal{B}_k^{\text{gC},\tau}(N_p)$ under restriction (using unique continuation). Consider the line bundle $E^{\text{agC},\tau}$ over $\mathcal{B}_k^{\text{gC},\tau}(N_p)$, pulled back from $E^{\text{agC},\sigma}$ under the map (157). From ζ^{agC} we obtain a section of $E^{\text{agC},\tau}$, and intersections of its zero set with the moduli spaces correspond to intersections of \mathcal{Z}^{agC} with those moduli spaces.

Observe that the pull-back of $E^{\text{agC},\tau}$ under $\Pi^{\text{[gC]},\tau}$ is exactly the bundle E^τ over $\mathcal{B}_k^\tau(N_p)$. Thus, we can pull back the section of $E^{\text{agC},\tau}$ and obtain a section of E^τ , which we can then take to be the one defining the U -maps in the τ model. We have a commutative diagram:

$$\begin{array}{ccc} M([x], [y]) & \longrightarrow & \mathcal{B}_k^\tau(N_p) \\ \Pi^{\text{[gC]},\tau} \downarrow & & \downarrow \Pi^{\text{[gC]},\tau} \\ M^{\text{agC}}([x], [y]) & \longrightarrow & \mathcal{B}_k^{\text{gC},\tau}(N_p), \end{array}$$

where the horizontal maps are given by restriction, and are one-to-one (by unique continuation). Since we have established that $\Pi^{\text{[gC]},\tau} : M([x], [y]) \rightarrow M^{\text{agC}}([x], [y])$ is a homeomorphism, we obtain an identification of the cut-down moduli spaces:

$$(158) \quad M([x], [y]) \cap \mathcal{Z}^\tau \cong M^{\text{agC}}([x], [y]) \cap \mathcal{Z}^{\text{agC}}$$

and

$$(159) \quad M^{\text{red}}([x], [y]) \cap \mathcal{Z}^\tau \cong M^{\text{agC},\text{red}}([x], [y]) \cap \mathcal{Z}^{\text{agC}}.$$

5.13. Orientations. In Section 4.6 we explained how the moduli spaces $M([x], [y])$ can be oriented using an orientation data set. That discussion can be adapted to global Coulomb gauge. To orient the spaces $M^{\text{agC}}([x], [y])$, we need to trivialize the determinant lines $\det(Q_\gamma^{\text{gC}})$. For arbitrary $x, y \in W^\sigma$, consider compact intervals $I = [t_1, t_2]$, paths $\gamma \in W^\tau(I \times Y)$ restricting to x and y on the two boundary components, and Fredholm operators of the form

$$(160) \quad P_\gamma^{\text{gC}} = (Q_\gamma^{\text{gC}}, -\Pi_1^{\text{gC},+}, \Pi_2^{\text{gC},-}) : \mathcal{T}_{1,\gamma}^{\text{gC},\tau}(I \times Y) \rightarrow \mathcal{V}_{0,\gamma}^{\text{gC},\tau} \oplus L^2(I; i\mathbb{R}) \oplus H_1^{\text{gC},+} \oplus H_2^{\text{gC},-}.$$

Here, Q_γ^{gC} has the same expression as in (138), and $\Pi_1^{\text{gC},+}$ is the composition of restriction to the boundary with spectral projection onto

$$H_1^{\text{gC},+} = \mathcal{K}_{1/2,x}^{\text{agC},+} \oplus i\mathbb{R} \subset \mathcal{T}_{1/2,x}^{\text{gC},\sigma}(Y) \oplus i\mathbb{R},$$

with $\mathcal{K}_{1/2,x}^{\text{agC},+} \subset \mathcal{K}_{1/2,x}^{\text{agC},\sigma}$ being the direct sum of all positive eigenspaces of the Hessian $\text{Hess}_{q,x}^{\bar{q},\sigma} - \epsilon$ for small, positive ϵ . The space $H_2^{\text{gC},-}$ and the projection $\Pi_2^{\text{gC},-}$ are defined

similarly, using the nonpositive eigenspaces of $\text{Hess}_{\mathbf{q},y}^{\tilde{g},\sigma} - \epsilon$. Note that since $\text{Hess}_{\mathbf{q},x}^{\tilde{g},\sigma}$ is invertible with real spectrum and is a compact perturbation of a self-adjoint and invertible operator, we see that $\mathcal{K}_{1/2,x}^{\text{agC},+} \oplus \mathcal{K}_{1/2,x}^{\text{agC},-} = \mathcal{K}_{1/2,x}^{\text{agC},\sigma}$ (cf. [KM07, p.313]).

We define an *orientation data set in Coulomb gauge* o^{gC} to consist of orientations $o_{[x],[y]}^{\text{gC}}$ for $\det(P_\gamma^{\text{gC}})$ (for any γ), satisfying the compatibility condition

$$o_{[x],[y]}^{\text{gC}} \cdot o_{[y],[z]}^{\text{gC}} = o_{[x],[z]}^{\text{gC}}.$$

An orientation data set in Coulomb gauge produces orientations on the moduli spaces $M^{\text{agC}}([x], [y])$ and $M^{\text{agC},\text{red}}([x], [y])$.

Proposition 5.45. *An orientation data set o (as in Section 4.6) naturally induces an orientation data set o^{gC} in Coulomb gauge, such that the homeomorphisms constructed in Proposition 5.36,*

$$M([x], [y]) \xrightarrow{\cong} M^{\text{agC}}([x], [y]),$$

are orientation-preserving, and so are the homeomorphisms

$$M^{\text{red}}([x], [y]) \xrightarrow{\cong} M^{\text{agC},\text{red}}([x], [y]).$$

Proof. Fix an orientation data set o . This trivializes the determinant lines $\det(P_\gamma)$, for the operators P_γ from (51). Let us focus on trajectories $\gamma \in W^\tau(I \times Y)$. We seek to trivialize the corresponding operators P_γ^{gC} . To go between P_γ and P_γ^{gC} , we follow the steps in the proof of Proposition 5.40, but considering operators defined on compact cylinders, and with spectral projections added at the boundary. Specifically, we can deform P_γ into an operator of the form

$$\widehat{P}_\gamma^{\text{gC}} = \begin{pmatrix} P_\gamma^{\text{gC}} & 0 \\ 0 & \widehat{J} \end{pmatrix},$$

where \widehat{J} is the analogue of \widehat{R} from (141), with spectral projections added at the boundary. One can check that \widehat{J} is bijective. Hence, a trivialization of $\det(P_\gamma)$ gives one of $\det(\widehat{P}_\gamma^{\text{gC}})$ and then one of $\det(P_\gamma^{\text{gC}})$. The resulting trivializations are compatible with concatenation, and hence combine into an orientation data set in Coulomb gauge.

The fact that the homeomorphisms are orientation-preserving is immediate from the construction. \square

Observe that Proposition 5.45 also implies that the identifications (158) and (159), between the cut-down moduli spaces, are orientation-preserving.

5.14. Monopole Floer homology in Coulomb gauge. The work done in this chapter allows us to rephrase the definition of monopole Floer homology in terms of configurations in Coulomb gauge.

We fix an admissible perturbation \mathbf{q} . The generators of $\widetilde{\mathcal{CM}}$ can be taken to be some of the stationary points $[x]$ of the vector field $\mathcal{X}_\mathbf{q}^{\text{agC},\sigma}$ on W^σ/S^1 ; precisely, those that are either in the interior of W^σ/S^1 , or on the boundary and stable. Indeed, by (118), these are in one-to-one correspondence with the generators $\mathcal{C}^o \cup \mathcal{C}^s$ considered in Section 4.7. Further, since the original generators are non-degenerate, so are the ones in Coulomb gauge, in the sense that $\ker(\mathcal{D}_x^\sigma \mathcal{X}_\mathbf{q}^{\text{agC},\sigma}) = 0$ at each stationary x ; see Lemma 5.33.

To define the differential on $\widetilde{\mathcal{CM}}$, we can use the moduli spaces $M^{\text{agC}}([x], [y])$ and $M^{\text{agC},\text{red}}([x], [y])$, consisting of trajectories of $\mathcal{X}_\mathbf{q}^{\text{agC},\sigma}$. By (119) and Proposition 5.36, these

are in one-to-one correspondence with the moduli spaces of monopoles considered in Section 4.7. Moreover, by Proposition 5.42, since the original moduli spaces are regular, so are the ones in Coulomb gauge. Specifically, this means that the operators $(\mathcal{D}_\gamma^\tau \mathcal{F}_q^{\text{gC}, \tau})|_{\mathcal{K}_{k, \gamma}^{\text{gC}, \tau}}$ (or, equivalently, Q_γ^{gC}) are surjective for all $[\gamma] \in M^{\text{agC}}([x], [y])$, except in the boundary-obstructed case, where the cokernel has dimension 1.

As shown in (153), we can define the relative gradings between generators $[x]$ and $[y]$ in Coulomb gauge by the index of the operator Q_γ^{gC} , where γ is a path from x to y in W^σ and that these are the same relative gradings as in \widetilde{CM} .

Moreover, we can orient the moduli spaces $M^{\text{agC}}([x], [y])$ and $M^{\text{agC}, \text{red}}([x], [y])$ using an orientation data set in Coulomb gauge, as in Section 5.13; see Proposition 5.45.

With this in mind, we define the differential $\check{\partial}$ on \widetilde{CM} by the same formulas as (52), (53), and (54), but using the moduli spaces $M^{\text{agC}}([x], [y])$, $M^{\text{agC}, \text{red}}([x], [y])$, instead of $M([x], [y])$, $M^{\text{red}}([x], [y])$.

Finally, the $\mathbb{Z}[U]$ -module structure on \widetilde{CM} can also be described in Coulomb gauge. We apply the equivalences (158) and (159) established in Section 5.12. Thus, the U -map is given by the formulas similar to (55), (56), and (57). For the new formulas, we instead intersect the moduli spaces $M^{\text{agC}}([x], [y])$, $M^{\text{agC}, \text{red}}([x], [y])$ with the zero set \mathcal{Z}^{agC} of a generic section of the complex line bundle $E^{\text{agC}, \sigma}$ over W^σ/S^1 .

6. FINITE-DIMENSIONAL APPROXIMATIONS WITH TAME PERTURBATIONS

6.1. Very compactness. Since monopole Floer homology is defined using a perturbation \mathfrak{q} , we must define an analogue of the Floer spectrum using this perturbation as well. In order to do this, we must recall a more general setting in which the spectrum can be defined.

Definition 6.1. *A smooth map $\eta : W \rightarrow W$ is called very compact if for all integers $k \geq 5$ and compact cylinders $Z = I \times Y$, the following two conditions are satisfied:*

(a) *The map η induces a continuous, functionally bounded map*

$$\hat{\eta} : W_{k-1}(Z) \rightarrow W_{k-1}(Z), \quad \text{and}$$

(b) *We can extend the differentials*

$$W \times \prod_{i=1}^m W \rightarrow W, \quad (x; v_1, \dots, v_m) \rightarrow (\mathcal{D}^m \eta)_x(v_1, \dots, v_m)$$

to continuous maps

$$W_k(Z) \times W_{k-1-i_1}(Z) \times \dots \times W_{k-1-i_m}(Z) \rightarrow W_{k-1-\sum i_s}(Z).$$

By the argument mentioned in Remark 4.6 (that is, working with configurations which are constant in the \mathbb{R} direction), we obtain the analogous regularity statements on Y as well. We call such a map η very compact because condition (a), applied with k instead of $k-1$, guarantees that the induced map $\eta : W_k \rightarrow W_k$ is functionally bounded, and therefore $\eta : W_k \rightarrow W_{k-1}$ is compact (in the sense that it takes bounded sets to precompact sets).

One example of a very compact map is the map $c : W \rightarrow W$ from (40).

Proposition 6.2. *Let $\eta : W \rightarrow W$ be a very compact map. Fix $k \geq 5$. Suppose that there exists a closed, bounded subset N of W_k such that the finite type trajectories of $l + \eta$ on W_k that are contained in N are actually contained in $W \cap U$ for an open set $U \subset N$. Then:*

- (i) *For $\lambda \gg 0$, trajectories of $l + p^\lambda \eta$ contained in N must be contained in U .*
- (ii) *We can define the Floer spectrum $\Sigma^{-W^{(-\lambda, 0)}} I^\lambda$ as in (3). This is independent of λ up to stable equivalence.*
- (iii) *Furthermore, if η is S^1 -equivariant, then the Floer spectrum can be constructed equivariantly, and is an invariant up to S^1 -equivariant stable equivalence.*

We omit the proof of Proposition 6.2, as it is analogous to that of [Man03, Theorem 1]. In particular, part (i) corresponds to [Man03, Proposition 3]. The results in [Man03] were for the special case where $\eta = c$, $N = \overline{B(2R)}$, and $U = B(R)$, as in Section 3. The properties of c used in those proofs were exactly those listed in Definition 6.1.

Remark 6.3. Very compactness was defined in [Man07, Definition 4]. There, condition (a) only required that $\eta : W_k \rightarrow W_{k-1}$ be a compact map. Proposition 5 in [Man07] claimed that Proposition 6.2 is true under this weaker hypothesis. However, this is in fact not sufficient. In [Man03], Step 1 in the proof of Proposition 3 requires functional boundedness of η to bound the derivatives of trajectories, and Step 3 requires continuity of $\hat{\eta}$ in order to do elliptic bootstrapping on $I \times Y$.

For future reference, we state a key lemma that is needed for the proof of Proposition 6.2(i). Its analogue is contained in the proof of Proposition 3 in [Man03].

Lemma 6.4. *Let $I \subseteq \mathbb{R}$ be a closed interval (possibly \mathbb{R}). Under the hypotheses of Proposition 6.2, suppose we have a sequence of eigenvalues $\lambda_n \rightarrow \infty$ and a sequence of trajectories*

$\gamma_n : I \rightarrow W$ of $l + p^{\lambda_n} \eta$, such that $\gamma_n(t) \in N$ for all $t \in I$. Then there exists a subsequence of γ_n for which the restrictions to any compact subinterval¹² $I' \Subset I$ converge in the C^∞ topology of $W(I' \times Y)$ to a trajectory of $l + \eta$.

The following will allow us to define a Floer spectrum which incorporates the perturbations used in defining monopole Floer homology.

Proposition 6.5. *Let \mathfrak{q} be a very tame perturbation. Then the map $\eta_{\mathfrak{q}} = \Pi_*^{\text{gC}} \mathfrak{q} : W \rightarrow W$ from (83) is very compact.*

Proof. Recall from Lemma 5.11 that $\eta_{\mathfrak{q}}$ is a controlled perturbation, in the sense of Definition 5.9.

Thus, to prove Property (a) in the definition of very compactness, we simply apply part (i) of Definition 5.9, with $k - 1$ instead of k . This says that $\eta_{\mathfrak{q}}$ extends to a continuous, functionally bounded map from $W_{k-1}(Z)$ to $W_{k-1}(Z)$.

For Property (b), without loss of generality, assume that i_m is the largest of the integers i_s . Let us study the regularity properties of the derivatives of \mathfrak{q} . We apply [KM07, Proposition 11.4.1(ii)], which states that for $p \geq 2$, the map

$$\mathcal{D}^m \hat{\mathfrak{q}} : (x, v_1, \dots, v_m) \mapsto \mathcal{D}_x^m \hat{\mathfrak{q}}(v_1, \dots, v_m),$$

extends to a continuous map from $\mathcal{C}_p(Z) \times \dots \times \mathcal{C}_p(Z) \times \mathcal{C}_j(Z)$ to $\mathcal{C}_j(Z)$, for $0 \leq j \leq p$. For clarity, this product is written to be compatible with the formula above, so the copy of $\mathcal{C}_j(Z)$ in the domain consists of tangent vectors and the left-most $\mathcal{C}_p(Z)$ consists of the point where we are computing the higher derivative. We let $j = k - 1 - i_m$ and

$$p = \begin{cases} k - 1 & \text{if } m = 1, \\ k - 1 - \max\{i_s \mid s \neq m\} & \text{if } m \geq 2. \end{cases}$$

Note that since $k - 1 \geq 4$, $k - 1 - \sum i_s \geq 0$, and $i_m = \max i_s$, we have that $k - 1 - i_s \geq 2$ for any $s \neq m$, so $p \geq 2$. Thus,

$$\mathcal{D}^m \hat{\mathfrak{q}} : \mathcal{C}_p(Z) \times \dots \times \mathcal{C}_p(Z) \times \mathcal{C}_{k-1-i_m}(Z) \rightarrow \mathcal{C}_{k-1-i_m}(Z)$$

is continuous. By pre-composing with the continuous inclusions of $\mathcal{C}_k(Z)$ and $\mathcal{C}_{k-1-i_s}(Z)$ into $\mathcal{C}_p(Z)$ for $s \neq m$ and post-composing with the continuous inclusion of $\mathcal{C}_{k-1-i_m}(Z)$ into $\mathcal{C}_{k-1-\sum i_s}(Z)$, we get that $\mathcal{D}^m \hat{\mathfrak{q}}$ extends to a map

$$\mathcal{C}_k(Z) \times \mathcal{C}_{k-1-i_1}(Z) \times \dots \times \mathcal{C}_{k-1-i_m}(Z) \rightarrow \mathcal{C}_{k-1-\sum i_s}(Z).$$

The requirement (b) in the definition of very compactness now follows from Lemma 3.1. \square

We want to apply Proposition 6.2 to l and

$$c_{\mathfrak{q}} := c + \eta_{\mathfrak{q}}.$$

Since both of the terms c and $\eta_{\mathfrak{q}}$ are very compact, so is $c_{\mathfrak{q}}$. However, observe that very compactness is only one of the requirements needed to do finite dimensional approximation. The others are smoothness and boundedness for flow trajectories, as in the statement of Proposition 6.2. (The containment in $W \cap U \subset W$ in the statement guarantees smoothness.) In our setting, we know from [KM07, Proposition 13.1.2 (i)] that for a tame perturbation, the perturbed Seiberg-Witten trajectories, and thus their global Coulomb projections, are smooth. With regard to boundedness, in Proposition 6.2, we want to take U to be the open ball of some radius $R \gg 0$ in W_k , and N to be the closed ball of radius $2R$. This is exactly

¹²For closed intervals $I, I' \subseteq \mathbb{R}$, we write $I' \Subset I$ if I' is compact and contained in the interior of I .

what was done for the unperturbed Seiberg-Witten trajectories in [Man03]. It follows from Proposition 1 of that paper that Seiberg-Witten trajectories inside $\overline{B(2R)}$ live inside $B(R)$, provided R was chosen large enough. The same proof works for perturbed Seiberg-Witten trajectories, since they satisfy similar compactness properties, as detailed in [KM07, Section 10.7].

Therefore, we are now able to define an analogue of the Floer spectrum with the finite-dimensional approximations of $l + c_q$ instead of $l + c$. We will denote this new spectrum by $\text{SWF}_q(Y, \mathfrak{s})$.

Proposition 6.6. *Let q be a very tame perturbation. Then $\text{SWF}_q(Y, \mathfrak{s})$ and $\text{SWF}(Y, \mathfrak{s})$ are S^1 -equivariantly stably homotopy equivalent.*

Proof. This is similar to the proof that $\text{SWF}(Y, \mathfrak{s})$ is independent of the Riemannian metric on Y ; see [Man03, Section 7]. The key ingredients are the fact that we can interpolate linearly between q and 0 such that the hypotheses of Proposition 6.2 are satisfied, and the fact that the homotopy type of the Conley index is invariant under perturbations. \square

Remark 6.7. It is worth comparing very compactness with the condition of being controlled, as in Definition 5.9. If q is very tame, both of these conditions are satisfied by η_q ; see Lemma 5.11 and Proposition 6.5. There is some overlap between the two conditions: for example, part (a) in the definition of very compactness is implied by part (i) in the definition of a controlled perturbation. However, neither condition is stronger than the other: part (b) in the definition of very compactness has no analog in the controlled condition; and part (ii) in the controlled condition has no analogue in very compactness due to the constraints on functional boundedness.

When working with the perturbation η_q , we will need to use both very compactness and the controlled condition. Very compactness is needed to do finite dimensional approximation, and the controlled compactness is necessary to study the properties of stationary points and trajectories in these approximations.

From now on, we always assume that our perturbation q is both very tame (Definition 4.9) and admissible (Definition 4.19). For the existence of such q , see Sections 4.4 and 4.5.

6.2. Strategy for the proof of Theorem 1.1. In order to relate monopole Floer homology to the spectrum, it suffices to instead work with $\text{SWF}_q(Y, \mathfrak{s})$ by Proposition 6.6. This latter invariant is clearly closer to monopole Floer homology due to the presence of the perturbation. However, we still need to relate the vector field $l + c_q$ used for monopole Floer homology to $l + p^\lambda c_q$ on W^λ . We will consider a vector field on W_k defined by taking finite-dimensional approximations of the non-linear part of \mathcal{X}_q :

$$(161) \quad \mathcal{X}_{q^\lambda}^{\text{gC}} := l + p^\lambda c_q = l + c + \eta_q^\lambda,$$

where $\eta_q^\lambda := p^\lambda c_q - c$. We will see that η_q^λ is very compact. Further, $\mathcal{X}_{q^\lambda}^{\text{gC}}$ induces a vector field $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ on the blow-up W^σ and thus a vector field $\mathcal{X}_q^{\text{agC}, \sigma}$ on the quotient W^σ/S^1 .

At this point we give an outline of how the proof of Theorem 1.1 is going to go.

For $N > 0$, let $\mathfrak{C}_{[-N, N]}$ be the set of stationary points of $\mathcal{X}_q^{\text{agC}, \sigma}$ of grading in $[-N, N]$. Their S^1 -orbits form sets of stationary points of $\mathcal{X}_q^{\text{gC}, \sigma}$ on W^σ . Let $\mathfrak{D}_{[-N, N]}$ denote the union of these orbits. We fix N sufficiently large so that the projection of $\mathfrak{D}_{[-N, N]}$ to the blow-down W contains all the stationary points of $\mathcal{X}_q^{\text{gC}}$; in particular, $\mathfrak{C}_{[-N, N]}$ should

contain all the irreducibles. Further, we assume that no reducible stationary point that is boundary-stable has grading less than $-N$. This ensures that the truncated chain complex $\widetilde{\mathcal{CM}}_{\leq N}(Y, \mathfrak{s}, \mathfrak{q})$ is generated by $\mathfrak{C}_{[-N, N]}$. Note that the homology of $\widetilde{\mathcal{CM}}_{\leq N}(Y, \mathfrak{s}, \mathfrak{q})$ agrees with $\widetilde{HM}(Y, \mathfrak{s}, \mathfrak{q})$ in degrees $\leq N - 1$.

Let

$$(162) \quad \mathcal{N} = \{x \in W_k^\sigma \mid d_{L_k^2}(x, \mathfrak{D}_{[-N, N]}) \leq 2\delta\},$$

where $d_{L_k^2}$ denotes L_k^2 distance, and $\delta > 0$ is chosen sufficiently small such that the only stationary points of $\mathcal{X}_q^{\text{agC}, \sigma}$ that are contained in \mathcal{N} are those in $\mathfrak{D}_{[-N, N]}$. Similarly, let

$$(163) \quad \mathcal{U} = \{x \in W_k^\sigma \mid d_{L_k^2}(x, \mathfrak{D}_{[-N, N]}) < \delta\} \subset \mathcal{N}.$$

Thus, \mathcal{N}/S^1 and \mathcal{U}/S^1 are closed, resp. open, neighborhoods of $\mathfrak{C}_{[-N, N]}$ in W_k^σ/S^1 .

Granted this, we will construct a chain complex \check{C}^λ determined by $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, and which will be identified with $\widetilde{\mathcal{CM}}_{\leq N}(Y, \mathfrak{s}, \mathfrak{q})$. The chain groups of \check{C}^λ will be generated by the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ that live in \mathcal{N}/S^1 (and hence in \mathcal{U}/S^1); in particular, this includes all irreducibles $[(a, s, \phi)]$ such that $(a, s\phi) \in \overline{B(2R)}$. We will see that these stationary points will necessarily be contained in the finite dimensional approximation $(W^\lambda)^\sigma/S^1$. The differential on \check{C}^λ will be defined analogously to monopole Floer homology (only counting trajectories that are contained entirely in $(W^\lambda)^\sigma/S^1$). That this will actually be a chain complex will come from a Morse-Smale stability condition—since we have non-degeneracy of the stationary points and regularity of the moduli spaces for $\mathcal{X}_q^{\text{agC}, \sigma}$, we will show that this holds for $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ as well, provided that λ is sufficiently large. Using the inverse function theorem, we will find a correspondence between the stationary points and isolated trajectories of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, on the one hand, and those of $\mathcal{X}_q^{\text{agC}, \sigma}$ on the other. This will give an explicit identification between \check{C}^λ and $\widetilde{\mathcal{CM}}_{\leq N}(Y, \mathfrak{s}, \mathfrak{q})$. Here we are using the work of Section 5, where we rephrased $\widetilde{\mathcal{CM}}$ in Coulomb gauge. By the setup, we will be able to relate the respective orientations of the moduli spaces, gradings, and U -actions as well.

However, \check{C}^λ can also be identified with a truncation of the Morse complex (for manifolds with S^1 actions) for $B(2R) \cap W^\lambda$ as in Section 2.8. It follows from Equation (22) that the homology of this Morse complex is isomorphic to $\widetilde{H}_{\leq M}^{S^1}(\text{SWF}_q(Y, \mathfrak{s}))$, for some $M > 0$. We can assume that $M > N$, and we get that

$$\widetilde{HM}_{\leq N-1}(Y, \mathfrak{s}, \mathfrak{q}) \cong \widetilde{H}_{\leq N-1}^{S^1}(\text{SWF}_q(Y, \mathfrak{s})).$$

By letting N tend to infinity and applying Proposition 6.6, we obtain the desired isomorphism in Theorem 1.1.

With the above strategy in mind, it will suffice to do most of our analysis (i.e. compactness and non-degeneracy of stationary points and trajectories) in W^σ rather than in the quotient W^σ/S^1 . We are able to do so due to the compactness of the residual gauge group S^1 .

6.3. Finite-dimensional approximations of perturbations. In order to carry out the strategy mentioned above, we need to understand the properties of the perturbation

$$\eta_q^\lambda = p^\lambda c_q - c = (p^\lambda c - c) + p^\lambda \eta_q,$$

and in particular of the term $p^\lambda \eta_q$. For notation, given a sequence $\lambda_n \rightarrow \infty$, we write π_n for p^{λ_n} . Recall that Θ_j is a constant such that $\|p^\lambda x\|_{L_j^2} \leq \Theta_j \|x\|_{L_j^2}$ for all $x \in W_j$ and we take $\Theta_0 = 1$.

Fact 6.8. (a) For all λ , the smoothed projection p^λ extends to a continuous, linear map $p^\lambda : W_j \rightarrow W_j$ such that $\|p^\lambda\| \leq \Theta_j$ for all j . (Here, $\|p^\lambda\|$ denotes the norm of p^λ as an operator on W_j .)

(b) If $\lambda_n \rightarrow \infty$, then $\pi_n = p^{\lambda_n} \rightarrow 1$ in the strong operator topology on W_j for all j .

Lemma 6.9. If q is a very tame perturbation, then $p^\lambda \eta_q$ is a controlled Coulomb perturbation for all λ .

Proof. Lemma 5.11 says that η_q is controlled. The claim now follows from Fact 6.8(a). \square

Lemma 6.10. The maps c_q , $p^\lambda \eta_q$, $p^\lambda c_q$ and η_q^λ are all very compact.

Proof. This follows from the very compactness of c , together with the very compactness of η_q (cf. Proposition 6.5), and Fact 6.8(a). \square

We also need the following result about c_q , which is not subsumed in very compactness.

Lemma 6.11. Let $Z = I \times Y$ for compact I . For $k \geq 3$ and $0 \leq j \leq k$, the map

$$\mathcal{D}c_q : W_k(Z) \rightarrow \text{Hom}(W_j(Z), W_j(Z))$$

is in C_{fb}^∞ . Therefore, extending by four-dimensional gauge, we have that

$$\mathcal{D}c_q : \mathcal{C}_k^{\text{gC}}(Z) \rightarrow \text{Hom}(\mathcal{T}_j^{\text{gC}}(Z), \mathcal{V}_j^{\text{gC}}(Z))$$

is in C_{fb}^∞ .

Proof. We write $c_q = c + \eta_q$. Since η_q is a controlled Coulomb perturbation by Lemma 5.11, we have that $\mathcal{D}\eta_q \in C_{\text{fb}}^\infty(W_k(Z), \text{Hom}(W_j(Z), W_j(Z)))$. Therefore, it suffices to show that $\mathcal{D}c$ is in C_{fb}^∞ as well. Write $x = (a, \phi)$.

By (40), we see that c is the composition of Π_*^{gC} with the vector field

$$(164) \quad X : W_k(Z) \rightarrow \mathcal{C}_k(Z), \quad (a, \phi) \mapsto (\tau(\phi, \phi), \rho(a) \cdot \phi).$$

We can explicitly compute

$$\begin{aligned} (\mathcal{D}_{(a, \phi)} X)(b, \psi) &= (\tau(\phi, \psi) + \tau(\psi, \phi), \rho(b) \cdot \phi + \rho(a) \cdot \psi) \\ (\mathcal{D}_{(a, \phi)}^2 X)((\alpha, \zeta), (b, \psi)) &= (\tau(\zeta, \psi) + \tau(\psi, \zeta), \rho(b) \cdot \zeta + \rho(\alpha) \cdot \psi) \\ (\mathcal{D}_{(a, \phi)}^3 X) &\equiv 0. \end{aligned}$$

From this, it is straightforward to apply the Sobolev multiplication $L_k^2 \times L_j^2 \rightarrow L_j^2$ to see that X satisfies the conditions of Lemma 5.12 with $n = \infty$. Therefore, we see that $\mathcal{D}c$ is in C_{fb}^∞ . \square

The vector field $\mathcal{X}_{q^\lambda}^{\text{gC}} = \mathcal{X}^{\text{gC}} + \eta_q^\lambda$ from (161) induces a vector field $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ on the blow-up. Similar to (91), this is given by the formula

$$\begin{aligned} \mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}(a, s, \phi) &= (\mathcal{X}_{q^\lambda}^0(a, s\phi), \Lambda_{q^\lambda}(a, s, \phi)s, D\phi + \widetilde{p^\lambda c_q}^1(a, s, \phi) - \Lambda_{q^\lambda}(a, s, \phi)\phi) \\ (165) \quad &= \left(*da + (p^\lambda c_q)^0(a, s\phi), \Lambda_{q^\lambda}(a, s, \phi)s, \right. \\ &\quad \left. D\phi + \int_0^1 \mathcal{D}_{(a, st\phi)}(p^\lambda c_q)^1(0, \phi)dt - \Lambda_{q^\lambda}(a, s, \phi)\phi \right), \end{aligned}$$

where

$$(166) \quad \Lambda_{q^\lambda} = \operatorname{Re} \langle \phi, (\widetilde{\mathcal{X}_{q^\lambda}^{\text{gC}}})^1(a, s, \phi) \rangle_{L^2}.$$

For $(a, s, \phi) \in W^\sigma$, we call $\Lambda_{q^\lambda}(a, s, \phi)$ the λ -spinorial energy. As in the case of the spinorial energy, the λ -spinorial energy of an irreducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ is zero. Analogous to the definition of $\mathcal{F}_q^{\text{gC}, \tau}$ in (92), we can use $\mathcal{G}^{\text{gC}}(Z)$ -equivariance to extend the equations defining the flow in (165) to define a section

$$(167) \quad \mathcal{F}_{q^\lambda}^{\text{gC}, \tau} : \widetilde{\mathcal{C}}^{\text{gC}, \tau}(Z) \rightarrow \mathcal{V}^{\text{gC}, \tau}(Z).$$

In the rest of this section we will describe several results about the dynamics of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$. These will be put to use, for instance, in Sections 7, 9, 12 and 13, where we will relate stationary points and trajectories of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ to those of $\mathcal{X}_q^{\text{gC}, \sigma}$.

Note that being a reducible stationary point $(a, 0)$ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is equivalent to solving

$$*da + (p^\lambda \eta_q)^0(a, 0) = 0, \quad (p^\lambda \eta_q)^1(a, 0) = 0,$$

since $p^\lambda c(a, 0) = 0$.

A reducible stationary point $(a, 0, \phi)$ of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ is determined by a reducible stationary point $(a, 0)$ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ and an L^2 -unit length eigenvector ϕ of the linear operator $D_{q^\lambda, a}$, where

$$(168) \quad D_{q^\lambda, a}(\phi) = \mathcal{D}_{(a, 0)}(\mathcal{X}_{q^\lambda}^{\text{gC}})^1(0, \phi) = D\phi + (p^\lambda)^1(\mathcal{D}_{(a, 0)}c_q(0, \phi)).$$

If $(a, 0, \phi)$ is a reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$, then the λ -spinorial energy is simply the eigenvalue of $D_{q^\lambda, a}$ corresponding to ϕ .

We now establish some important analytic properties of the perturbed equations.

Lemma 6.12. *Fix $k \geq 3$.*

(a) *For a bounded subset $K \subset W_k$ and a bounded subset $J \subset W_j$ (with $0 \leq j \leq k$), the set*

$$\{p^\lambda(\mathcal{D}_{(a, \phi)}c_q(b, \psi)) \mid (a, \phi) \in K, (b, \psi) \in J, \lambda \geq 0\}$$

is bounded in L_j^2 . These uniform bounds also hold if we include $\lambda = \infty$.

(b) *For a bounded subset $K^\sigma \subset W_k^\sigma \subset (\ker d^*)_k \oplus \mathbb{R} \oplus L_k^2(Y; \mathbb{S})$ and a bounded subset $J^\sigma \subset \mathcal{T}_j^{\text{gC}, \sigma}|_{K^\sigma}$ (with $0 \leq j \leq k$), the set*

$$\{\mathcal{D}_x^\sigma(p^\lambda c_q)^\sigma(v) \mid x \in K^\sigma, (x, v) \in J^\sigma, \lambda \geq 0\}$$

is bounded in L_j^2 . These uniform bounds also hold if we include $\lambda = \infty$.

(c) *If $x_n \rightarrow x$ in W_k^σ and $\lambda_n \rightarrow \lambda$ (possibly ∞), then $\widetilde{p^{\lambda_n} c_q}^1(x_n) \rightarrow \widetilde{p^\lambda c_q}^1(x)$ in $L_k^2(Y; \mathbb{S})$.*

Furthermore, $\widetilde{\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}}^1(x_n) \rightarrow \widetilde{\mathcal{X}_{q^\lambda}^{\text{gC}}}^1(x)$ in $L_{k-1}^2(Y; \mathbb{S})$. In this case, we also have $\Lambda_{q^{\lambda_n}}(x_n) \rightarrow \Lambda_{q^\lambda}(x)$ and thus $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}, \sigma}(x_n) \rightarrow \mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}(x)$ in $\mathcal{T}_{k-1}^{\text{gC}, \sigma}$.

(d) *Same as (c), but for four-dimensional configurations on compact cylinders.*

Proof. (a) We have $c_q = c + \eta_q$. In the proof of Lemma 6.11, it was shown that $\mathcal{D}c_q \in C_{\text{fb}}^0(W_k, \operatorname{Hom}(W_j, W_j))$. We thus have L_j^2 bounds on the set

$$\{\mathcal{D}_{(a, \phi)}c_q(b, \psi) \mid (a, \phi) \in K, (b, \psi) \in J\}.$$

Fact 6.8(a) now gives uniform L_j^2 bounds on

$$\{p^\lambda(\mathcal{D}_{(a,\phi)}c_q(b,\psi)) \mid (a,\phi) \in K, (b,\psi) \in J, \lambda \geq 0\}.$$

(b) Recall that $\mathcal{D}^\sigma(p^\lambda c_q)^\sigma$ is computed in two steps. First, we differentiate $(p^\lambda c_q)^\sigma$, thought of as a map to $L^2(Y; i\mathbb{R}) \times \mathbb{R} \times L^2(Y; \mathbb{S})$, as opposed to a section of $\mathcal{T}_k^{\text{gC}}$. Then, we apply L^2 orthogonal projection to $\mathcal{T}_k^{\text{gC},\sigma}$. We also have the analogous construction when extending to L_j^2 completions. Since $\Pi_{\mathcal{T}_{j,(a,s,\phi)}^{\text{gC},\sigma}}(b,r,\psi) = (b,r,\psi - \text{Re}\langle\psi,\phi\rangle_{L^2}\phi)$, we see that $\Pi_{\mathcal{T}_j^{\text{gC},\sigma}}$ is L_j^2 bounded in terms of the L_k^2 norm of (a,s,ϕ) and L_j^2 norm of (b,r,ψ) . Therefore, we focus on the boundedness of the derivative of $(p^\lambda c_q)^\sigma$ as a map to the extended space $L_k^2(Y; i\mathbb{R}) \times \mathbb{R} \times L_k^2(Y; \mathbb{S})$. From the definitions we compute

(169)

$$\begin{aligned} \mathcal{D}_{(a,s,\phi)}(p^\lambda c_q)^\sigma(b,r,\psi) &= \left(\mathcal{D}_{(a,s,\phi)}(p^\lambda c_q)^0(b,r\phi + s\psi), \right. \\ &\quad \text{Re}\langle \widetilde{p^\lambda c_q^1}(a,s,\phi), \phi \rangle_{L^2} r + \text{Re}\langle \widetilde{p^\lambda c_q^1}(a,s,\phi), \psi \rangle_{L^2} s + \text{Re}\langle \mathcal{D}_{(a,s,\phi)} \widetilde{p^\lambda c_q^1}(b,r,\psi), \phi \rangle_{L^2} s, \\ &\quad \left. \text{Re}\langle \mathcal{D}_{(a,s,\phi)} \widetilde{p^\lambda c_q^1}(b,r,\psi), \phi \rangle_{L^2} \phi + \text{Re}\langle \widetilde{p^\lambda c_q^1}(a,s,\phi), \psi \rangle_{L^2} \phi + \text{Re}\langle \widetilde{p^\lambda c_q^1}(a,s,\phi), \phi \rangle_{L^2} \psi \right), \end{aligned}$$

where $\widetilde{p^\lambda c_q^1}(a,s,\phi) = \int_0^1 \mathcal{D}_{(a,sr\phi)}(p^\lambda c_q)^1(0,\phi) dr$. It follows from Lemma 6.11 and Fact 6.8(a) that

$$\widetilde{p^\lambda c_q^1}(a,s,\phi), \mathcal{D}_{(a,s,\phi)} \widetilde{p^\lambda c_q^1}(b,r,\psi)$$

are L_j^2 bounded in terms of the L_k^2 bounds on (a,s,ϕ) and the L_j^2 bounds on (b,r,ψ) . Sobolev multiplication again provides the desired bounds on (169).

(c) A smooth section of W_k to the L_j^2 completion of TW_k with $j \leq k$ induces a smooth vector field on the blow-up with the same regularity. By Proposition 6.5 (which applies for any Sobolev coefficient at least three), we see that if $x_n \rightarrow x$ in W_k^σ , we have $\widetilde{c_q^1}(x_n) \rightarrow \widetilde{c_q^1}(x)$ in L_k^2 . Applying Fact 6.8 gives $\widetilde{p^{\lambda_n} c_q}(x_n) \rightarrow \widetilde{p^\lambda c_q}(x)$ in L_k^2 as well. Since $\mathcal{X}_{q^\lambda}^{\text{gC}} = l + p^\lambda c_q$ and $l : W_k \rightarrow W_{k-1}$ are smooth for $k \geq 2$, we obtain the desired convergence results by similar arguments.

(d) The argument is similar to the one for (c). □

It turns out that we can give more explicit bounds on $(p^\lambda c_q)^1$:

Lemma 6.13. *There exists a constant $C_{k,j,M}$, depending only on $0 \leq j \leq k$ and M , such that if $x = (a,\phi) \in W_k$ satisfies $\|x\|_{L_k^2} \leq M$ then $\|(p^\lambda c_q)^1(x)\|_{L_j^2} \leq C_{k,j,M} \|\phi\|_{L_j^2}$ for any λ .*

Note that $\|\phi\|_{L_j^2} \leq \|x\|_{L_k^2} \leq M$, so Lemma 6.13 implies that $\|(p^\lambda c_q)^1(x)\|_{L_j^2}$ is bounded above uniformly by a constant $C'_{k,j,M}$. We already knew this, owing to the very compactness of $p^\lambda c_q$. The new content of Lemma 6.13 is that it gives a stronger (linear) bound when $\|\phi\|_{L_j^2}$ is small.

Proof of Lemma 6.13. Since $\|p^\lambda\| \leq \Theta_j$, it suffices to find a constant C such that $\|c_q^1(x)\|_{L_j^2} \leq C \|\phi\|_{L_j^2}$. First, we bound c^1 . Recall that

$$c^1(a,\phi) = \rho(a) \cdot \phi - Gd^* \tau(\phi, \phi) \phi.$$

Since a is L_k^2 -bounded, Sobolev multiplication implies

$$\|\rho(a) \cdot \phi\|_{L_j^2} \leq C_1 \|\phi\|_{L_j^2}.$$

Because ϕ is L_k^2 -bounded, so is $Gd^*\tau(\phi, \phi)$. Again, by Sobolev multiplication, we obtain

$$\|Gd^*\tau(\phi, \phi)\phi\|_{L_j^2} \leq C_2 \|\phi\|_{L_j^2}.$$

Therefore, it remains to find bounds on $\eta_q^1(a, \phi)$. We have

$$\eta_q^1(a, \phi) = \mathbf{q}^1(a, \phi) - Gd^*\mathbf{q}^0(a, \phi)\phi.$$

Since \mathbf{q} is a very tame perturbation, we obtain L_k^2 bounds on $Gd^*\mathbf{q}^0(a, \phi)$ and thus linear bounds on the second term. Therefore, it remains to give linear bounds on $\|\mathbf{q}^1(a, \phi)\|_{L_j^2}$ in terms of $\|\phi\|_{L_j^2}$. We may write $\mathbf{q}^1(a, \phi) = \int_0^1 \mathcal{D}_{(a, r\phi)} \mathbf{q}^1(0, \phi) dr$, because the S^1 -equivariance of \mathbf{q}^1 implies $\mathbf{q}^1(a, 0) = 0$. Since the $(a, r\phi)$ are L_k^2 bounded and $\mathcal{D}_{(a, r\phi)} \mathbf{q}^1$ is linear in ϕ , the very tameness of \mathbf{q} implies linear L_j^2 -bounds on $\mathcal{D}_{(a, r\phi)} \mathbf{q}^1(0, \phi)$, depending only on $\|\phi\|_{L_j^2}$. Integrating gives linear bounds on $\|\mathbf{q}^1(a, \phi)\|_{L_j^2}$ in terms of $\|\phi\|_{L_j^2}$. This now completes the proof. \square

In order to make further connection with the finite-dimensional setting of the spectrum, we must ensure that the trajectories of $\mathcal{X}_{q^\lambda}^{\mathbf{g}^C, \sigma}$ are actually contained in the finite-dimensional blow-up of W^λ , assuming control over the λ -spinorial energy. We begin with the non-blown-up case.

Lemma 6.14. *If $\gamma(t)$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\mathbf{g}^C}$ contained in $B(2R)$, then $\gamma(t)$ is contained in W^λ . In particular, $\gamma(t)$ is smooth for each t .*

Proof. Let $\gamma(t) = (a(t), \phi(t))$. Fix κ with $|\kappa| \geq \lambda$ and let $\gamma_\kappa(t)$ be the L^2 projection of $\gamma(t)$ to the κ -eigenspace of l . It suffices to show that $\gamma_\kappa(t) = 0$. Since $\gamma(t)$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\mathbf{g}^C, \tau}$, we have

$$-\frac{d}{dt}\gamma(t) = l(\gamma(t)) + (p^\lambda c_q)(\gamma(t)).$$

Since $|\kappa| \geq \lambda$, we see $-\frac{d}{dt}\gamma_\kappa(t) = l(\gamma_\kappa(t)) = \kappa\gamma_\kappa(t)$. Thus, $\gamma_\kappa(t) = e^{-\kappa t}\gamma_\kappa(0)$. Suppose that $\gamma_\kappa(t) \not\equiv 0$. Therefore, $\gamma_\kappa(t)$, and thus $\gamma(t)$, are unbounded in the L_k^2 -norm on Y as $|t|$ increases. This contradicts $\gamma(t) \in B(2R)$ for all t . Therefore, $\gamma_\kappa(t) \equiv 0$ and $\gamma(t)$ is contained in W^λ . \square

We seek an analogous result to Lemma 6.14, but in the blow-up. Apart from the $2R$ bound on the projections of trajectories in the blow-down, we will also need to assume a bound ω on the absolute values of spinorial energies Λ_q . With N , \mathcal{N} and \mathcal{U} as in Section 6.2, note that Λ_q is bounded on the compact set $\mathfrak{D}_{[-N, N]}$, and hence it is bounded on its neighborhood \mathcal{N} . The boundedness can be explained as follows. Since \mathbf{q} is controlled and our choice of k guarantees that $k - 1 \geq 2$, we have that Λ_q is continuous on W_{k-1}^σ . Therefore, Λ_q is bounded on \mathcal{N} since it is a compact subset of W_{k-1}^σ . We shall choose $\omega > 0$ so that any point $x \in \mathcal{N}$ has

$$(170) \quad |\Lambda_q(x)| < \omega.$$

Lemma 6.15. *For sufficiently large λ , the following is true. If $\gamma(t) = (a(t), s(t), \phi(t))$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ contained in $B(2R)^\sigma$ and the λ -spinorial energy of $\gamma(t)$ is in $[-\omega, \omega]$ for all t , then $(a(t), \phi(t)) \in W^\lambda$ for all $t \in \mathbb{R}$. In particular, $(a(t), \phi(t))$ is smooth for all $t \in \mathbb{R}$.*

Proof. Fix $\lambda \gg \omega$. Fix κ with $|\kappa| \geq \lambda$ and let $(a_\kappa(t), \phi_\kappa(t))$ be the projection of $(a(t), \phi(t))$ to the κ -eigenspace of l . Since $(a(t), s(t)\phi(t))$ is contained in W^λ by Lemma 6.14, we have that $a_\kappa(t) \equiv 0$. It suffices to show that $\phi_\kappa(t) \equiv 0$. Let $\nu(t) = \Lambda_{q^\lambda}(a(t), s(t), \phi(t))$. Since $\gamma(t)$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$, we have by (165) that

$$\begin{aligned} -\frac{d}{dt}\phi(t) &= D\phi(t) + \widetilde{p^\lambda c_q}^1(a(t), s(t), \phi(t)) - \nu(t)\phi(t) \\ &= D\phi(t) + \int_0^1 \mathcal{D}_{(a(t), rs(t)\phi(t))}(p^\lambda c_q)^1(0, \phi(t))dr - \nu(t)\phi(t). \end{aligned}$$

Since $p^\lambda c_q(x) \in W^\lambda$ for all $x \in W_k$, we have that the projection of $\widetilde{p^\lambda c_q}^1(a(t), s(t), \phi(t))$ to the κ -eigenspace of D is trivial. Therefore, we have

$$-\frac{d}{dt}\phi_\kappa(t) = D\phi_\kappa(t) - \nu(t)\phi_\kappa(t) = (\kappa - \nu(t))\phi_\kappa(t).$$

By assumption, $\kappa - \nu(t) \gg 0$ for all t . If $\phi_\kappa(t) \not\equiv 0$, then $\|\phi_\kappa(t)\|_{L^2(Y)}$ is unbounded. However, $\|\phi_\kappa(t)\|_{L^2(Y)} \leq \|\phi(t)\|_{L^2(Y)} \equiv 1$, which is thus a contradiction. \square

For future reference, we will write $L_k^2(Y)$ (resp. $L_k^2(Z)$) to refer to the L_k^2 -norm of any relevant object, e.g. a connection, spinor, etc., on Y (resp. Z). If $Z = I \times Y$, we note that $L_{j+1}^2(Z)$ -convergence implies pointwise $L_j^2(Y)$ -convergence, uniform on compact sets of I . We will also write x, a, s, ϕ (without t) if we want to treat this object as a section of a bundle over $I \times Y$ instead of as a path of sections over Y .

7. STATIONARY POINTS

Throughout Sections 7–13 we fix the following:

- a very tame, admissible perturbation \mathfrak{q} ; this means that \mathfrak{q} is very tame in the sense of Definition 4.9, the stationary points of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ are non-degenerate, and the associated moduli spaces of trajectories are regular; we will impose additional conditions on \mathfrak{q} in Proposition 7.11 and in Proposition 10.2;
- a Sobolev index $k \geq 5$;
- a bound $R > 0$ such that all the stationary points and finite type trajectories of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$ are contained in $B(2R) \subset W_k$;
- a value $N > 0$ specifying a grading range $[-N, N]$, a closed neighborhood \mathcal{N} and an open neighborhood $\mathcal{U} \subset \mathcal{N}$ of the set of stationary points of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ in that grading range, as in Section 6.2; we also assume that the projection of \mathcal{N} to the blow-down is contained in $B(2R)$ and that N is chosen large enough to contain each reducible stationary point $(a, 0, \phi)$ of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ where ϕ is an eigenvector of $D_{\mathfrak{q},a}$ with smallest positive eigenvalue;
- a strict bound ω on the absolute values of the spinorial energies of points in \mathcal{N} , as in (170).

In analogy with the finite-dimensional setting, for $M \subset W$, we will use M^σ to denote the closure in W^σ of the preimage of $W - (\ker d^* \oplus 0)$ under the projection from W^σ to W ; and similarly for Sobolev completions. In particular, we can talk about $B(2R)^\sigma$. Note that \mathcal{N} is a subset of $B(2R)^\sigma$.

Recall that it is our goal to identify some of the stationary points and trajectories of $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$ with those of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$, for λ sufficiently large. In this section, we deal with stationary points. First, in Section 7.1 we show that the stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$ contained in \mathcal{N}/S^1 are close to stationary points of $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$ in \mathcal{N}/S^1 , for λ sufficiently large. Then, in Section 7.2 an inverse function theorem argument shows that inside \mathcal{N}/S^1 , the stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$ are in one-to-one correspondence with those of $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$. In Section 7.3 we prove that the nearby approximate stationary points are non-degenerate. (This will be needed later, when we define a Morse complex for $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$.) In Section 7.4, we study stationary points outside of \mathcal{N}/S^1 . We rephrase these results with gradings of stationary points in Section 9.1.

As before, if we have a sequence $\lambda_n \rightarrow \infty$, we write π_n to denote p^{λ_n} .

7.1. Convergence. We first point out a convergence result for stationary points in W (not in the blow-up).

Lemma 7.1. *Suppose that x_n is a sequence of stationary points for $\mathcal{X}_{\mathfrak{q}^{\lambda_n}}^{\text{gC}}$ in $B(2R)$ where $\lambda_n \rightarrow \infty$. Then, there is a subsequence that converges in W_k to x , a stationary point for $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$. Further, if x_n are reducible, so is x .*

Proof. This follows from applying Lemma 6.4 to constant trajectories, noting that $c_{\mathfrak{q}}$ is a very compact map by Proposition 6.5. If $x_n = (a_n, 0) \in B(2R)$ is a sequence of reducibles, then clearly the limit must be of the form $x = (a, 0)$. \square

We now need an analogous compactness result on the blow-up. It turns out that very compactness of $c_{\mathfrak{q}}$ is not sufficient; we will use some of the other properties of controlled Coulomb perturbations to do this.

Lemma 7.2. *Fix $\epsilon > 0$. There exists $b \gg 0$ such that for all $\lambda > b$ the following is true. If $x \in \mathcal{N} \subset W^\sigma$ is a zero of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$, then there exists $x' \in \mathcal{N}$ such that $\mathcal{X}_q^{\text{gC},\sigma}(x') = 0$ and x, x' have L_k^2 -distance at most ϵ in $L_k^2(Y; iT^*Y) \oplus \mathbb{R} \oplus L_k^2(Y; \mathbb{S})$.*

Proof. Suppose this is not true. Then, we can find a sequence $\lambda_n \rightarrow \infty$ and corresponding zeros $x_n = (a_n, s_n, \phi_n)$ of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC},\sigma}$ in \mathcal{N} , none of which are within L_k^2 -distance ϵ of a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$ in \mathcal{N} . We will contradict this by finding a subsequence converging to such a stationary point.

Since the x_n are in \mathcal{N} , which is L_k^2 -bounded, we can extract a subsequence that converges to some $x = (a, s, \phi)$ in L_{k-1}^2 . Further, after passing to another subsequence, we can assume the x_n are either all reducible or all irreducible.

First, suppose the x_n are reducible, that is, $s_n = 0$. By Lemma 7.1, the convergence $a_n \rightarrow a$ is in the stronger L_k^2 norm. Moreover, $a \in (\ker d^*)_k$ is such that $(a, 0)$ is a stationary point of $\mathcal{X}_q^{\text{gC}}$. Recall that ϕ_n are eigenvectors for $D_{q^{\lambda_n}, a_n}$ with L^2 -norm equal to 1. We claim that ϕ is an eigenvector of $D_{q,a} = D + \mathcal{D}_{(a,0)}(c_q)^1(0, \cdot)$, and that the convergence $\phi_n \rightarrow \phi$ is also in L_k^2 .

Let κ_n be the associated eigenvalues for ϕ_n , i.e., the λ_n -spinorial energies of x_n . Since $x_n \rightarrow x$ in L_{k-1}^2 , by Lemma 6.12 (c) we have $|\Lambda_q(x)| \leq \omega$. By applying Lemma 6.12 (c) again, this time to the λ_n -spinorial energies of x_n , we see that $\Lambda_{q^{\lambda_n}}(x_n) \rightarrow \Lambda_q(x)$. Therefore, we also have a bound $\omega' \geq \omega$ on $|\Lambda_{q^{\lambda_n}}(x_n)| = |\kappa_n|$. After passing to a subsequence, we can assume that the κ_n converge to some value κ . Then, we have

$$\begin{aligned}
 \kappa_n \phi_n &= D_{q^{\lambda_n}, a_n}(\phi_n) \\
 &= D\phi_n + \mathcal{D}_{(a_n, 0)}(\pi_n c_q)^1(0, \phi_n) \\
 &= D\phi_n + \pi_n(\mathcal{D}_{(a_n, 0)}c_q(0, \phi_n))^1.
 \end{aligned}
 \tag{171}$$

Now, since $(a_n, 0)$ converges to $(a, 0)$ in L_k^2 and $(0, \phi_n)$ converge to $(0, \phi)$ in L_{k-1}^2 , the very compactness of c_q guarantees that $\mathcal{D}_{(a_n, 0)}c_q(0, \phi_n)$ converges to $\mathcal{D}_{(a, 0)}c_q(0, \phi)$ in L_{k-1}^2 . As $\pi_n \rightarrow 1$ in the strong operator topology on W_{k-1} , we must also have

$$\pi_n(\mathcal{D}_{(a_n, 0)}c_q(0, \phi_n))^1 \rightarrow (\mathcal{D}_{(a, 0)}c_q(0, \phi))^1 \text{ in } L_{k-1}^2.$$

On the other hand, we have

$$\kappa_n \phi_n \rightarrow \kappa \phi \text{ in } L_{k-1}^2.$$

Finally, we have that $D\phi_n \rightarrow D\phi$ in L_{k-2}^2 , because $\phi_n \rightarrow \phi$ in L_{k-1}^2 . Since the convergence of $\pi_n \mathcal{D}_{(a_n, 0)}(c_q)^1(0, \phi_n)$ and $\kappa_n \phi_n$ is in L_{k-1}^2 , we must in fact have that $D\phi_n \rightarrow D\phi$ in L_{k-1}^2 by (171). Thus, ϕ_n converges to ϕ in L_k^2 and $\kappa_n \phi_n = D_{q^{\lambda_n}, a_n}(\phi_n) \rightarrow D_{q,a}(\phi) = \kappa \phi$ in L_k^2 . Thus, ϕ is an eigenvalue of $D_{q,a}$, so $x = (a, 0, \phi)$ is a reducible stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$. Since \mathcal{N} is closed in the L_k^2 norm and the convergence $x_n \rightarrow x$ is in L_k^2 , we get that $x \in \mathcal{N}$, providing the contradiction.

We now assume that our sequence (a_n, s_n, ϕ_n) consists of irreducibles. If $s_n \geq \delta > 0$ for all n , Lemma 7.1 guarantees that (a_n, s_n, ϕ_n) will converge in W_k^σ to an irreducible stationary point (a, s, ϕ) of $\mathcal{X}_q^{\text{gC},\sigma}$, since $\mathcal{X}_{q^{\lambda_n}}^{\text{gC},\sigma}$ is conjugate to $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ via the blow-down. This is again a contradiction.

The final case is when (a_n, s_n, ϕ_n) is a sequence of irreducible stationary points in $B(2R)^\sigma$ with $s_n \rightarrow 0$. By Lemma 7.1, in the blow-down we have $(a_n, s_n \phi_n) \rightarrow (a, 0)$ in L_k^2 . Since

$x_n \in \mathcal{N}$, we can find an upper bound on $\|\phi_n\|_{L_k^2}$. By Lemma 6.13, we have a constant C such that

$$|s_n| \cdot \|D\phi_n\|_{L_k^2} = \|D(s_n\phi_n)\|_{L_k^2} = \|(\pi_n c_q)^1(a_n, s_n\phi_n)\|_{L_k^2} \leq C|s_n|.$$

Thus, we obtain L_{k+1}^2 -bounds on ϕ_n . After passing to a subsequence, we get that (a_n, s_n, ϕ_n) converges in L_k^2 to $(a, 0, \phi)$. By Lemma 6.12 (c), we have that $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}, \sigma}(a_n, s_n, \phi_n) \rightarrow \mathcal{X}_q^{\text{gC}, \sigma}(a, 0, \phi)$, and thus $(a, 0, \phi)$ is a reducible stationary point. Moreover, since $\Lambda_{q^{\lambda_n}}(a_n, s_n, \phi_n) = 0$, by taking the limit $n \rightarrow \infty$ we obtain $\Lambda_q(a, 0, \phi) = 0$, i.e. ϕ is in the kernel of $D_{q,a}$. This is impossible, because the non-degeneracy condition for reducibles requires that 0 is not in the spectrum of $D_{q,a}$. \square

Remark 7.3. The proof of Lemma 7.2 shows that if the zero x of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ is reducible, then we can choose the nearby zero x' of $\mathcal{X}_q^{\text{gC}, \sigma}$ to be reducible as well. Furthermore, since $\Lambda_{q^\lambda}(x)$ is close to $\Lambda_q(x')$, we see that x' can be chosen to be stable (resp. unstable) when x is stable (resp. unstable).

Corollary 7.4. *For $\lambda \gg 0$, if x is a zero of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ in \mathcal{N} , then x is actually in \mathcal{U} and $|\Lambda_{q^\lambda}(x)| < \omega$.*

Proof. We get that $x \in \mathcal{U}$ by choosing ϵ sufficiently small in Lemma 7.2. If we had a sequence $\lambda_n \rightarrow \infty$ and corresponding $x_n \in \mathcal{N}$ with $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}, \sigma}(x_n) = 0$ and $|\Lambda_{q^{\lambda_n}}(x_n)| \geq \omega$, then after extracting a subsequence, and making use of Lemma 7.2, in the limit we would get a zero of $\mathcal{X}_q^{\text{gC}, \sigma}$ inside \mathcal{N} with $|\Lambda_q| \geq \omega$. This is a contradiction. \square

Corollary 7.5. *For $\lambda \gg 0$, all the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ in \mathcal{N} live inside the finite-dimensional blow-up $(W^\lambda)^\sigma$.*

Proof. This follows from Corollary 7.4 and Lemma 6.15. \square

Observe that the results of this subsection also apply to the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ and $\mathcal{X}_q^{\text{agC}, \sigma}$, in the quotient W^σ/S^1 .

7.2. Stability. In the previous section we showed that, for λ large, the zeros of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ in \mathcal{N} are close to the zeros of $\mathcal{X}_q^{\text{agC}, \sigma}$. Our next goal is to show that near each stationary point of $\mathcal{X}_q^{\text{agC}, \sigma}$ there is exactly one stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, again for λ large. The proof will be an application of the implicit function theorem, using the fact that we have chosen the perturbation q so that the zeros of $\mathcal{X}_q^{\text{agC}, \sigma}$ are nondegenerate.

Recall that p^λ is (roughly) the orthogonal projection to W^λ , modified so that it becomes a smooth function of λ for $\lambda \in (0, \infty)$. In order to be able to apply the implicit function theorem at $\lambda = \infty$, we need to find a suitable identification of the interval $(0, \infty]$ with $[0, 1)$, so that after this identification we get differentiability at zero.

For the present subsection, let us index (with multiplicity) the eigenvalues of l by $(\lambda_n)_{n \geq 0}$ by the condition that $|\lambda_n| \leq |\lambda_{n+1}|$ for all n ; note that $\lim_{n \rightarrow \infty} |\lambda_n| = \infty$. Recall that we write

$$\mathcal{X}_q^{\text{gC}} = l + c_q.$$

Pick a homeomorphism $f : (0, \infty] \rightarrow [0, 1)$ with the following properties:

- The restriction of f to $(0, \infty)$ is a strictly decreasing diffeomorphism onto $(0, 1)$;

- $\lim_{n \rightarrow \infty} |\lambda_n|^2 f(|\lambda_{n+1}|) = \infty$.

The second property means that f does not decrease too fast near infinity. We can achieve it, for example, by requiring that $f(|\lambda_{n+1}|) = \frac{1}{|\lambda_n|} + \frac{1}{n}$ for large n .

Lemma 7.6. *The map*

$$h : W_k \times (-1, 1) \rightarrow W_{k-1}, \quad h(x, r) = x - p^{f^{-1}(|r|)}(x)$$

is continuously differentiable, with $Dh_{(x,0)}(0, 1) = 0$ for all x .

Proof. Continuous differentiability away from $r = 0$ is standard, taking into account that p^λ are smoothed projections as in (42).

Since $h(x, 0) = 0$, it suffices to show that

$$\lim_{r \rightarrow 0} \frac{h(x, r)}{r} = 0,$$

or, equivalently

$$(172) \quad \lim_{\lambda \rightarrow \infty} \frac{\|x - p^\lambda(x)\|_{L_{k-1}^2}}{f(\lambda)} = 0.$$

The eigenspaces of l are orthogonal in L^2 . Pick an L^2 -orthonormal sequence of eigenvectors w_n for λ_n . Let us write

$$x = \sum_n x_n w_n,$$

for a sequence of real numbers $(x_n)_{n \geq 0}$. We have $\|x\|_{L^2}^2 = \sum_n |x_n|^2$.

Note that

$$l^{k-1}(x) = \sum_n \lambda_n^{k-1} x_n w_n.$$

Since the L_{k-1}^2 norm is equivalent to the one defined using l as a differential operator instead of ∇ , we can write

$$\|x\|_{L_{k-1}^2}^2 \approx \sum_n \sum_{j=0}^{k-1} |\lambda_n|^{2j} |x_n|^2 \approx \sum_n |\lambda_n|^{2k-2} |x_n|^2.$$

In fact, we have $x \in W_k$, so

$$\|x\|_{L_k^2}^2 \approx \sum_n |\lambda_n|^{2k} |x_n|^2 < \infty.$$

We have

$$\|x - p^{|\lambda_n|}(x)\|_{L_{k-1}^2} \approx \sum_{m>n} |\lambda_m|^{2k-2} |x_m|^2 < \frac{1}{|\lambda_n|^2} \sum_{m>n} |\lambda_m|^{2k} |x_m|^2 \leq C \frac{\|x\|_{L_k^2}^2}{|\lambda_n|^2}$$

for some constant C independent of x and λ_n .

Recall that we chose f so that $\lim_{n \rightarrow \infty} |\lambda_n|^2 f(|\lambda_{n+1}|) = \infty$. From here we get:

$$\lim_{n \rightarrow \infty} \frac{\|x - p^{|\lambda_n|}(x)\|_{L_{k-1}^2}}{f(|\lambda_{n+1}|)} = 0.$$

The claim (172) follows: For any $\lambda \gg 0$, we choose n such that $\lambda \in [|\lambda_n|, |\lambda_{n+1}|]$, and then we use the fact that both the numerator and the denominator of the right hand side of (172) are nonincreasing functions of λ . \square

Let $[x_0] \in W^\sigma/S^1$ be a non-degenerate, irreducible zero of $\mathcal{X}_q^{\text{agC},\sigma}$. By Lemma 5.33, non-degeneracy means that the linearization

$$\mathcal{D}_{[x_0]}^\sigma(\mathcal{X}_q^{\text{agC},\sigma}) : \mathcal{K}_{k,[x_0]}^{\text{agC},\sigma} \rightarrow \mathcal{K}_{k-1,[x_0]}^{\text{agC},\sigma}$$

is an invertible linear operator.

Proposition 7.7. *Let $[x] \in W_k^\sigma/S^1$ be a non-degenerate stationary point of $\mathcal{X}_q^{\text{agC},\sigma}$. Then, for any sufficiently small neighborhood $U_{[x]}$ of $[x]$ in W_k^σ/S^1 , for $\lambda \gg 0$ (depending on $U_{[x]}$) there is a unique $[x_\lambda] \in U_{[x]}$ satisfying $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}([x_\lambda]) = 0$.*

Proof. Consider the vector field $S : (W_k^\sigma/S^1) \times (-1, 1) \rightarrow \mathcal{K}_{k-1}^{\text{agC},\sigma} \times \mathbb{R}$ given by

$$S([x], r) = (\mathcal{X}_{q^{f^{-1}(|r|)}}^{\text{agC},\sigma}([x]), r).$$

If we choose a representative $x \in W_k^\sigma$ for $[x]$, we can write

$$S([x], r) = ([(\Pi^{\text{agC},\sigma} \circ (l + p^{f^{-1}(|r|)} c_q)^\sigma(x))], r).$$

Using Lemma 5.23 and its adaption to $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ together with Lemma 7.6, we deduce that S is continuously differentiable. Furthermore, from Lemma 7.6, the derivative of S at $([x], 0)$ is the same as the one obtained without the factor $p^{f^{-1}(|r|)}$; that is, it equals $\begin{pmatrix} \mathcal{D}_{[x]}^\sigma(\mathcal{X}_q^{\text{agC},\sigma}) & 0 \\ 0 & 1 \end{pmatrix}$, which is invertible by hypothesis.

Note that $(W_k^\sigma/S^1) \times (-1, 1)$ is a Banach manifold with boundary. Let $\mathcal{D}(W_k^\sigma/S^1 \times (-1, 1))$ be its double, equipped with the associated involution ι . Since S is tangent to the boundary of $(W_k^\sigma/S^1) \times (-1, 1)$, S extends to the double in an ι -invariant way. We now apply the inverse function theorem to S (on the double) at $([x], 0)$. We get that there is a unique solution of $S([y], r) = (0, r)$ near $([x], 0)$, for small $r > 0$. Define $\lambda = f(r)$, and write $[x_\lambda] = [y]$.

Note that if $[x]$ was irreducible (i.e. in the interior of $(W_k^\sigma/S^1) \times (-1, 1)$), then the nearby solution, $[x_\lambda]$, is an irreducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$. If $[x]$ was reducible (i.e. on the boundary of $(W_k^\sigma/S^1) \times (-1, 1)$), by ι -invariance, the nearby $[x_\lambda]$ is also on the boundary (i.e., it is a reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$). \square

Recall that in Section 6.2 we denoted by $\mathfrak{C}_{[-N,N]}$ the set of stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$ with grading in the interval $[-N, N]$ (or, equivalently, those inside \mathcal{N}/S^1). For simplicity, we write $\mathfrak{C}_\mathcal{N}$ for $\mathfrak{C}_{[-N,N]}$. Moreover, we denote by $\mathfrak{C}_\mathcal{N}^\lambda$ the set of stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ that live in \mathcal{N}/S^1 .

Combining the results of this subsection and the previous one, we have the following:

Corollary 7.8. *For $\lambda \gg 0$, there is a one-to-one correspondence*

$$\Xi_\lambda : \mathfrak{C}_\mathcal{N}^\lambda \rightarrow \mathfrak{C}_\mathcal{N}.$$

This correspondence preserves the type of stationary point (irreducible, stable, unstable).

Proof. There are only finitely many stationary points in $\mathfrak{C}_\mathcal{N}$. Thus, we can find $\epsilon > 0$ such that the neighborhoods chosen in Proposition 7.7 around each $[x] \in \mathfrak{C}_\mathcal{N}$ are the balls of L_k^2 radius ϵ around those points. We let $\Xi_\lambda^{-1}([x])$ be the unique zeros of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ inside those

balls. Since the points in $\mathfrak{C}_{\mathcal{N}}$ are actually inside the smaller open set $\mathcal{U}/S^1 \subset \mathcal{N}/S^1$, we can assume that $\Xi_{\lambda}^{-1}([x])$ are in \mathcal{N}/S^1 as well. By Lemma 7.2, we get that any stationary point of $\mathcal{X}_{q^{\lambda}}^{\text{agC},\sigma}$ from \mathcal{N}/S^1 must be of the form $\Xi_{\lambda}^{-1}([x])$ for some $[x] \in \mathfrak{C}_{\mathcal{N}}$. Thus, Ξ_{λ} is a one-to-one correspondence.

The claim about preserving type follows from Remark 7.3 and the arguments at the end of the proof of Proposition 7.7. \square

From now on, we will generally use $[x_{\infty}]$ to denote a stationary point of $\mathcal{X}_q^{\text{agC},\sigma}$ (corresponding to $\lambda = \infty$). If we fix $[x_{\infty}] \in \mathfrak{C}_{\mathcal{N}}$, observe that $\Xi_{\lambda}^{-1}([x_{\infty}])$, which is a stationary point of $\mathcal{X}_{q^{\lambda}}^{\text{agC},\sigma}$, is smooth as a function of λ ; this is guaranteed by the implicit function theorem. We will write $[x_{\lambda}]$ for $\Xi_{\lambda}^{-1}([x_{\infty}])$.

Finally, note that since the reducible stationary points of $\mathcal{X}_q^{\text{gC},\sigma}$ all correspond to non-zero eigenvalues, the same holds for the corresponding reducible stationary points of $\mathcal{X}_{q^{\lambda}}^{\text{agC},\sigma}$, assuming $\lambda \gg 0$.

It turns out we also can obtain analogous results for reducible stationary points in the blow-down. This will be useful when analyzing the reducible stationary points which are in $(B(2R) \cap W^{\lambda})^{\sigma}$, but not in \mathcal{N} ; see Proposition 9.10 below.

Lemma 7.9. *Fix $\epsilon > 0$. For $\lambda \gg 0$, there is a one-to-one correspondence in $B(2R)$ between reducible stationary points x_{∞} of $\mathcal{X}_q^{\text{gC}}$ and reducible stationary points x_{λ} of $\mathcal{X}_{q^{\lambda}}^{\text{gC}}$; further, x_{λ} is ϵ -close to x .*

Proof. It is straightforward to verify that since the stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$ are non-degenerate, so are the reducible stationary points of $\mathcal{X}_q^{\text{gC}}$ when restricting to the reducible locus in W_k . By the implicit function theorem, as used in the proof of Proposition 7.7, we obtain that near each reducible stationary point x_{∞} of $\mathcal{X}_q^{\text{gC}}$, there is a unique nearby reducible stationary point x_{λ} of $\mathcal{X}_{q^{\lambda}}^{\text{gC}}$. (This is in fact easier than Proposition 7.7, since there is no need to blow-up or quotient by S^1 .) On the other hand, by Lemma 7.1, reducible stationary points of $\mathcal{X}_{q^{\lambda}}^{\text{gC}}$ are necessarily nearby to reducible stationary points of $\mathcal{X}_q^{\text{gC}}$. This establishes the desired correspondence. Finally, Lemma 6.14 implies that x_{λ} is also in W^{λ} . \square

7.3. Hyperbolicity. Recall that we chose q such that any stationary point x of $\mathcal{X}_q^{\text{gC},\sigma}$ is non-degenerate, that is, the Hessian $\text{Hess}_{q,x}^{\tilde{g},\sigma}$ is invertible. By Lemma 5.22 (b), $\text{Hess}_{q,x}^{\tilde{g},\sigma}$ has real spectrum, and hence is a hyperbolic operator (i.e., the spectrum of its complexification is disjoint from the real axis).

For the approximate vector field $\mathcal{X}_{q^{\lambda}}^{\text{gC},\sigma}$, we define the \tilde{g} -Hessian in the blow-up by analogy with (103):

$$(173) \quad \text{Hess}_{q^{\lambda},x}^{\tilde{g},\sigma} = \Pi_x^{\text{agC},\sigma} \circ \mathcal{D}_x^{\tilde{g},\sigma} \mathcal{X}_{q^{\lambda}}^{\text{gC},\sigma} : \mathcal{K}_{k,x}^{\text{agC},\sigma} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC},\sigma}.$$

As before, if x is a stationary point of $\mathcal{X}_{q^{\lambda}}^{\text{gC},\sigma}$, then it does not matter which connection we use to differentiate $\mathcal{X}_{q^{\lambda}}^{\text{gC},\sigma}$ at x . Therefore, we can simply write

$$\text{Hess}_{q^{\lambda},x}^{\tilde{g},\sigma} = \Pi_x^{\text{agC},\sigma} \circ \mathcal{D}_x^{\sigma} \mathcal{X}_{q^{\lambda}}^{\text{gC},\sigma}.$$

We will say that x is a *hyperbolic* stationary point if $\text{Hess}_{\mathfrak{q}^\lambda, x}^{\tilde{g}, \sigma}$ is hyperbolic. (Generally, it will no longer be the case that $\text{Hess}_{\mathfrak{q}^\lambda, x}^{\tilde{g}, \sigma}$ has real spectrum.)

If x is in \mathcal{N} , then from Corollary 7.5 we know that if x lives in the finite dimensional blow-up $(W^\lambda)^\sigma$. Furthermore, since $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}, \sigma} = l + p^\lambda c_{\mathfrak{q}}$ maps W^λ to W^λ , then x being hyperbolic (as above, in infinite dimensions) implies that $[x]$ is also hyperbolic as a stationary point of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC}, \sigma}$ restricted to the finite dimensional approximation $(W^\lambda)^\sigma/S^1$. Proving hyperbolicity for these stationary points will be the first step towards defining Morse homology with $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC}, \sigma}$ on $(W^\lambda)^\sigma/S^1$ as in Section 2.8.

Let us start by analyzing the Hessian on the blow-up more carefully.

Writing $x = (a, s, \phi) \in W_k^\sigma$, we have

$$(174) \quad l^\sigma(a, s, \phi) = (*da, \langle D\phi, \phi \rangle_{L^2} s, D\phi - \langle D\phi, \phi \rangle_{L^2} D\phi).$$

Thus, we obtain for $v = (b, r, \psi) \in \mathcal{T}_{j,x}^{\text{gC}, \sigma}$

$$(175) \quad \begin{aligned} \mathcal{D}_x^\sigma l^\sigma(v) &= (*db, \langle D\phi, \phi \rangle_{L^2} r + 2 \operatorname{Re} \langle D\psi, \phi \rangle_{L^2} s, \Pi_\phi^\perp(D\psi - \langle D\phi, \phi \rangle_{L^2} \psi - 2 \operatorname{Re} \langle D\psi, \phi \rangle_{L^2} \phi)) \\ &= (*db, \langle D\phi, \phi \rangle_{L^2} r + 2 \operatorname{Re} \langle D\psi, \phi \rangle_{L^2} s, \Pi_\phi^\perp(D\psi) - \langle D\phi, \phi \rangle_{L^2} \psi) \\ &= (*db, \langle D\phi, \phi \rangle_{L^2} r + 2 \operatorname{Re} \langle D\psi, \phi \rangle_{L^2} s, D\psi - \langle D\psi, \phi \rangle_{L^2} \phi - \langle D\phi, \phi \rangle_{L^2} \psi) \end{aligned}$$

where Π_ϕ^\perp denotes L^2 orthogonal projection onto $\{\psi' \in L_{j-1}^2(Y; \mathbb{S}) \mid \operatorname{Re} \langle \phi, \psi' \rangle_{L^2} = 0\}$. Here we are using that $\Pi_\phi^\perp(\psi) = \psi$ since $(b, r, \psi) \in \mathcal{T}_{j,x}^\sigma$. Again, note that here we are taking derivatives with respect to the L^2 metric. Observe that L_k^2 bounds on (a, s, ϕ) and L_j^2 bounds on (b, r, ψ) easily give bounds on the r -component of (175).

Recall from Lemma 5.22 that $\text{Hess}_{\mathfrak{q}, x}^{\tilde{g}, \sigma}$ has Fredholm index 0, regardless of whether x is a stationary point. Further, we have that $\text{Hess}_{\mathfrak{q}^\lambda, x}^{\tilde{g}, \sigma}$ is a compact perturbation of $\text{Hess}_{\mathfrak{q}, x}^{\tilde{g}, \sigma}$ (roughly since both $p^\lambda c_{\mathfrak{q}}$ and $c_{\mathfrak{q}}$ are compact as maps from W_k to W_{k-1}), and thus has Fredholm index 0. Further, since the inclusion $\mathcal{K}_{k,x}^{\text{agC}, \sigma} \rightarrow \mathcal{K}_{k-1,x}^{\text{agC}, \sigma}$ is compact, we see that for any $z \in \mathbb{C}$, after complexifying,

$$(176) \quad \operatorname{ind}(\text{Hess}_{\mathfrak{q}, x}^{\tilde{g}, \sigma} - zI) = 0.$$

Proposition 7.10. *For λ sufficiently large, the stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}, \sigma}$ inside \mathcal{N} are hyperbolic. As a consequence, among the stationary points of the restriction of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC}, \sigma}$ to the finite dimensional space $(W^\lambda)^\sigma/S^1$, all those in \mathcal{N}/S^1 are hyperbolic (and hence non-degenerate).*

Proof. Recall from Section 3.4 that we can use an equivalent L_k^2 metric on W (and hence W^σ), where we use the operator l in place of the covariant derivatives. The statement is independent of which metric we use, so we opt for l in this proof.

Suppose the claim is false. Consider a sequence $x_n = (a_n, s_n, \phi_n) \in \mathcal{N}$ of stationary points of $\mathcal{X}_{\mathfrak{q}^{\lambda_n}}^{\text{gC}, \sigma}$ which are non-hyperbolic, where $\lambda_n \rightarrow \infty$. Note that we have L_k^2 -bounds on each component of x_n . Moreover, by Lemma 7.2, there exists a subsequence of x_n which converges in W_k^σ to x , a stationary point of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}, \sigma}$. By our assumption on \mathfrak{q} , the stationary point x is hyperbolic. Since each x_n is non-hyperbolic, we may find a sequence of real numbers κ_n such that, after complexifying, $\text{Hess}_{\mathfrak{q}^{\lambda_n}, x_n}^{\tilde{g}, \sigma} - i\kappa_n I$ is not invertible. Since $\operatorname{ind}(\text{Hess}_{\mathfrak{q}^{\lambda_n}, x_n}^{\tilde{g}, \sigma} - i\kappa_n I) = 0$ by (176), these operators have non-trivial kernel, and thus we

may find non-zero $v_n \in \mathcal{K}_{k,x_n}^{\text{agC},\sigma} \otimes \mathbb{C}$ such that $\text{Hess}_{q^{\lambda_n},x_n}^{\tilde{g},\sigma}(v_n) = i\kappa_n v_n$. For the rest of the proof, we drop the complexified notation for simplicity.

After rescaling, we can assume that $\|v_n\|_{L_k^2} = 1$. We will show that there exists a subsequence which converges in $\mathcal{K}_k^{\text{agC},\sigma}$ to a non-trivial element v of $\mathcal{K}_{k,x}^{\text{agC},\sigma}$ for which $\text{Hess}_{q,x}^{\tilde{g},\sigma}(v) = i\kappa v$, for some $\kappa \in \mathbb{R}$. This will contradict the hyperbolicity of x .

We write $v_n = (b_n, r_n, \psi_n)$. Our first step is to prove that the κ_n are bounded, as are the L_{k+1}^2 -norms of v_n .

By Lemma 6.12(b), we get a uniform bound on the L_k^2 norms of $\Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma (p^{\lambda_n} c_q)^\sigma(v_n)$. Further, since x_n is a stationary point of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC},\sigma}$, we can write

$$\begin{aligned} (177) \quad i\kappa_n v_n &= \text{Hess}_{q^{\lambda_n},x_n}^{\tilde{g},\sigma}(v_n) \\ &= \Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma \mathcal{X}_{q^{\lambda_n}}^{\text{gC},\sigma}(v_n) \\ &= \Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma l^\sigma(v_n) + \Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma (p^{\lambda_n} c_q)^\sigma(v_n). \end{aligned}$$

By definition, $\Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma l^\sigma(v_n) = \mathcal{D}_{x_n}^\sigma l^\sigma(v_n) - \alpha_n(0, 0, i\phi_n)$ for some sequence of real numbers α_n . The L_k^2 -bounds on x_n and v_n give L_{k-1}^2 -bounds on $\mathcal{D}_{x_n}^\sigma l^\sigma(v_n)$ by (175). By continuity, $\Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma l^\sigma(v_n)$ is L_{k-1}^2 -bounded. Because $\|\phi_n\|_{L_{k-1}^2} \geq 1$, the sequence α_n is bounded. Thus, $\alpha_n(0, 0, i\phi_n)$ is L_k^2 -bounded.

Now note that $\langle D\psi_n, \phi_n \rangle_{L^2} \phi_n$ is L_k^2 bounded. By (175), we then have that $\mathcal{D}_{x_n}^\sigma l^\sigma(v_n) - (*db_n, 0, D\psi_n - \langle D\phi_n, \phi_n \rangle_{L^2} \psi_n)$ is L_k^2 bounded; here we are using the observation that the L_k^2 bounds on x_n and v_n guarantee bounds on the r -component of $\mathcal{D}_{x_n}^\sigma l^\sigma(v_n)$. Combining this with the above discussion, we have that

$$\Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma l^\sigma(v_n) + \Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma (p^{\lambda_n} c_q)^\sigma(v_n) - (*db_n, 0, D\psi_n - \langle D\phi_n, \phi_n \rangle_{L^2} \psi_n)$$

is L_k^2 -bounded. We now see (177) shows that

$$(178) \quad i\kappa_n(b_n, r_n, \psi_n) - (*db_n, r_n, D\psi_n - \langle D\phi_n, \phi_n \rangle_{L^2} \psi_n)$$

is L_k^2 -bounded. However, note that the operator $(b, r, \psi) \mapsto (*db, r, D\psi - \langle D\phi_n, \phi_n \rangle_{L^2} \psi)$ is L^2 self-adjoint. It follows from here that $i\kappa_n v_n$ and $(*db_n, r_n, D\psi_n - \langle D\phi_n, \phi_n \rangle_{L^2} \psi_n)$ are (real) orthogonal in the complexification, with respect to the L^2 inner product. Given that we are working with the L_k^2 inner products that are defined using l as the derivative, we see that the orthogonality also holds with respect to L_k^2 . Thus, the L_k^2 bounds on the quantity in (178) imply that $\kappa_n v_n$ is L_k^2 -bounded, and thus the κ_n are bounded. On the other hand, we also obtain L_k^2 -bounds on $(*db_n, r_n, D\psi_n)$ and thus the v_n are bounded in L_{k+1}^2 by the ellipticity of $*d$ and D .

Thus, we can find a subsequence of the v_n that converges in $\mathcal{K}_k^{\text{agC},\sigma}$ to an element v , which necessarily has $\|v\|_{L_k^2} = 1$. Because $\mathcal{K}_k^{\text{agC},\sigma}$ is closed in the tangent bundle to W_k^σ , we have $v \in \mathcal{K}_k^{\text{agC},\sigma}$ as well. After passing to a further subsequence, $\kappa_n \rightarrow \kappa$, for some $\kappa \in \mathbb{R}$. We have

$$\text{Hess}_{q^{\lambda_n},x_n}^{\tilde{g},\sigma}(v_n) - \text{Hess}_{q,x_n}^{\tilde{g},\sigma}(v_n) = \Pi^{\text{agC},\sigma} \circ \mathcal{D}_{x_n}^\sigma (p^{\lambda_n} c_q - c_q)^\sigma(v_n) \rightarrow 0 \text{ in } L_{k-1}^2.$$

By assumption, the sequence $\text{Hess}_{q^{\lambda_n},x_n}^{\tilde{g},\sigma}(v_n)$ is the sequence $i\kappa_n v_n$. Moreover, $\text{Hess}_{q,x_n}^{\tilde{g},\sigma}(v_n)$ converges to $\text{Hess}_{q,x}^{\tilde{g},\sigma}(v)$ in L_{k-1}^2 by the continuity of the Hessian (cf. Lemma 5.23). Thus,

there exists a non-zero $v \in \mathcal{K}_k^{\text{agC},\sigma}$ such that $\text{Hess}_{\tilde{q},x}^{\tilde{g}}(v) = i\kappa v$. This contradicts the hyperbolicity of x . \square

7.4. Other stationary points. Proposition 7.10 was only about the stationary points in \mathcal{N} . Recall that \mathcal{N} is a subset of $B(2R)^\sigma \subset W_k^\sigma$. We do not have any control over the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ outside $B(2R)^\sigma$, but we can say a bit more about the ones in $B(2R)^\sigma$ (and not necessarily in \mathcal{N}).

First, recall from Corollary 7.5 that, for $\lambda \gg 0$, all the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ in \mathcal{N} are actually inside the finite-dimensional blow-up $(W^\lambda)^\sigma$.

Second, by applying Proposition 6.2 to $N = \overline{B(2R)}$ and U being the blow-down of \mathcal{U} , we see that for $\lambda \gg 0$, all the irreducible stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ in $B(2R)^\sigma$ are actually in $\mathcal{U} \subset \mathcal{N}$.

Some of the reducible solutions to $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ are in \mathcal{N} , and hence close to reducible zeros of $\mathcal{X}_q^{\text{gC},\sigma}$ with grading in $[-N, N]$. However, there will be other reducibles in $B(2R)^\sigma$ which may not be in \mathcal{N} . We now study these other reducibles. A reducible $(a, 0, \phi)$ has to satisfy:

$$- * da = (p^\lambda c_q)^0(a, 0)$$

and

$$D_{q^\lambda,a}(\phi) = D\phi + (p^\lambda)^1(\mathcal{D}_{(a,0)}c_q(0, \phi)) = \kappa\phi,$$

for some $\kappa \in \mathbb{R}$. Note that $(a, 0) \in W^\lambda$, but $(a, 0, \phi)$ may or may not be in $(W^\lambda)^\sigma$.

Observe that D , $(p^\lambda)^1$ and $\mathcal{D}_{(a,0)}c_q(0, \cdot)$ are all L^2 self-adjoint maps on spinors. (We are using here that the \tilde{g} metric agrees with the L^2 metric at reducibles.) Nevertheless, the product of $(p^\lambda)^1$ and $\mathcal{D}_{(a,0)}c_q(0, \cdot)$, and hence the operator $D_{q^\lambda,a}$, may not be self-adjoint. On the other hand, for the real numbers λ_i^\bullet defined in Section 3.4, the restriction of $D_{q^{\lambda_i^\bullet},a}$ to (the spinorial part of) W^λ is self-adjoint. This is because for $\lambda = \lambda_i^\bullet$, the map p^λ is the honest L^2 projection onto W^λ ; therefore, for $(a, 0, \phi) \in (W^\lambda)^\sigma$ we can write

$$D + (p^\lambda)^1\mathcal{D}_{(a,0)}c_q(0, \cdot) = D + (p^\lambda)^1\mathcal{D}_{(a,0)}c_q(0, \cdot)(p^\lambda)^1,$$

and the right hand side is self-adjoint.

We will focus on the reducible stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ that are in $(B(2R) \cap W^\lambda)^\sigma$ for $\lambda = \lambda_i^\bullet$. We then have the following strengthening of Proposition 7.10:

Proposition 7.11. *We can choose the admissible perturbation q such that for any $\lambda \in \{\lambda_1^\bullet, \lambda_2^\bullet, \dots\}$ sufficiently large, the restriction of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ to $(B(2R) \cap W^\lambda)^\sigma$ has only hyperbolic (and hence non-degenerate) stationary points.*

Proof. As part of the proof of existence of admissible perturbations, Kronheimer and Mrowka showed in [KM07, Section 12.6] that for a residual (and hence nonempty) set of tame perturbations q , the reducible stationary points of \mathcal{X}_q are non-degenerate. The key point is in the proof of Lemma 12.6.2 in [KM07]: There is a large enough space of tame (in fact, very tame) perturbations q^\perp (given by cylinder functions) that vanish at the reducible locus, such that in the tangent space to any q^\perp we can find a $\delta q^\perp = \text{grad } \delta f$, such that the Hessian of $\delta f|_V$ is any chosen S^1 -equivariant self-adjoint endomorphism of $V = \ker(D_{q^\perp,a})$.

We can adapt this proof to $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$, and pick q such that all the reducibles are non-degenerate for a given λ . Indeed, we now need to find a $\delta q^\perp = \text{grad } \delta f$ such that the \tilde{g} -Hessian of $\delta f|_V$ is any S^1 -equivariant self-adjoint endomorphism of $V = \ker(D_{(q^\perp)^\lambda,a}) \subset$

W^λ . Since we are at a reducible, on the spinorial part we have that the \tilde{g} -Hessian is the same as the usual Hessian. Further, because $\lambda = \lambda_i^\bullet$, we have that p^λ is the L^2 orthogonal projection to W^λ . Hence, we have

$$\text{Hess}(\delta f|_{W^\lambda}) = p^\lambda \circ \text{Hess}(\delta f)|_{W^\lambda},$$

and we can arrange so that this equals any chosen S^1 -equivariant self-adjoint endomorphism of the spinors in W^λ . From here we get the same freedom in choosing $\text{Hess}(\delta f|_V)$, for any subspace V of spinors in W^λ . The other arguments from [KM07, Section 12.6] can then be easily adapted to our setting.

Since λ is part of a countable collection $\{\lambda_n^\bullet\}$, we can find \mathfrak{q} such that the reducibles are non-degenerate for all such λ . Note that non-degeneracy implies hyperbolicity for reducibles, because the relevant operators are self-adjoint.

Since the irreducible stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$ in $B(2R)^\sigma$ are actually in $\mathcal{U} \subset \mathcal{N}$, by applying Proposition 7.10 we can arrange so that these irreducibles are hyperbolic as well. \square

Proposition 7.11 will be useful in showing that the vector field $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$ in $(B(2R) \cap W^\lambda)^\sigma$ is a Morse-Smale equivariant quasi-gradient. We can then construct a Morse homology group from it, and show that it is the same as Morse homology in $\mathcal{N} \cap (W^\lambda)^\sigma$, in a certain grading range $[-N, N]$. Indeed, we will show that all the other reducible points in $(B(2R) \cap W^\lambda)^\sigma$ cannot be in this grading range; see Proposition 9.10 below. Before we can discuss and define gradings on stationary points as in Section 2, we must first establish that $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}}$ is indeed a Morse quasi-gradient. This is the subject of the following section. We return to discuss gradings on stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$ in Section 9 and the Morse-Smale condition in Section 10.

8. THE APPROXIMATE FLOW AS A MORSE EQUIVARIANT QUASI-GRADIENT

Throughout this section we assume that the eigenvalue cut-off λ is of the form λ_i^\bullet for $i \gg 0$.

Note that $\mathcal{X}_q^{\text{gC}} = l + c_q$ is the gradient of the \mathcal{L}_q functional with respect to the \tilde{g} metric. However, the maps p^λ are defined in terms of projections with respect to the usual L^2 metric. As discussed in Remark 3.4, the vector field

$$\mathcal{X}_{q^\lambda}^{\text{gC}} = l + p^\lambda c_q$$

on W^λ is neither the L^2 nor the \tilde{g} gradient of the restriction of \mathcal{L}_q to W^λ . In fact, there is no reason for the derivative of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ at stationary points to have real spectrum (as it would happen for a gradient vector field, with respect to any metric).

Nevertheless, in this section we will be able to prove the following.

Proposition 8.1. *We can choose the admissible perturbation q such that for all $\lambda = \lambda_i^\bullet$ with $i \gg 0$, the vector field $\mathcal{X}_{q^\lambda}^{\text{gC}}$ on $W^\lambda \cap B(2R)$ is a Morse equivariant quasi-gradient, in the sense of Definition 2.18.*

As discussed in Section 2, having a Morse-Smale equivariant quasi-gradient suffices in order to construct (equivariant) Morse homology; the first step towards this is establishing the Morse condition (cf. Definition 2.18). The additional Morse-Smale condition on trajectories will be shown in Section 10.

In view of Lemma 2.19, to check that $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient we need three things:

- that the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ are hyperbolic;
- that the operators $\mathcal{D}_x \mathcal{X}_{q^\lambda}^{\text{gC}}$ at reducible stationary points x are self-adjoint;
- part (d) of Definition 2.18.

Hyperbolicity of the stationary points was already checked in Proposition 7.11. Self-adjointness of the operators $\mathcal{D}_x \mathcal{X}_{q^\lambda}^{\text{gC}} = l + p^\lambda \mathcal{D}_x c_q$ at reducibles follows from the fact that the metrics \tilde{g} and L^2 coincide there.

We are left to verify part (d) of Definition 2.18. Sections 8.1-8.6 below are devoted to proving this.

Proposition 8.2. *For each $\lambda \gg 0$, there exists a smooth function*

$$F_\lambda : W^\lambda \cap B(2R) \rightarrow \mathbb{R}$$

such that

$$(179) \quad \frac{1}{4} \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2 \leq dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \leq 4 \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2.$$

In particular, $dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \geq 0$, with equality only at the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$.

Remark 8.3. In the statement of Proposition 8.2, the constants $\frac{1}{4}$ and 4 are quite arbitrary. They could be replaced (at the expense of increasing λ) by $\frac{1}{C}$ and C , for any $C > 1$.

8.1. Control away from the stationary points. Since we know that $\mathcal{X}_q^{\text{gC}}$ is the \tilde{g} -gradient of the perturbed CSD functional \mathcal{L}_q , the first guess is to take F_λ to be \mathcal{L}_q . Then, the desired condition holds away from neighborhoods of the stationary points:

Lemma 8.4. *Fix $\epsilon > 0$. Then, for all $\lambda \gg 0$, we have*

$$(180) \quad \frac{1}{4} \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2 < d\mathcal{L}_q(\mathcal{X}_{q^\lambda}^{\text{gC}}) < 4 \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2$$

at any point in $W^\lambda \cap B(2R)$ which is at L_{k-1}^2 distance at least ϵ from all stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ in $B(2R)$.

Proof. We suppose this is not true. Then there exists sequences $\lambda_n \rightarrow \infty$ and $x_n \in W^{\lambda_n} \cap B(2R)$ such that x_n is L_{k-1}^2 distance at least ϵ from each stationary point of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ and $(d\mathcal{L}_q)_{x_n}(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}})$ violates (180). Without loss of generality, we assume that the first inequality in (180) is violated. (The case of the second inequality is similar.) Since the x_n are L_k^2 bounded, we can extract a subsequence which converges in L_{k-1}^2 to some element x . We see that

$$(d\mathcal{L}_q)_{x_n}(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}) \rightarrow (d\mathcal{L}_q)_x(\mathcal{X}_q^{\text{gC}}).$$

Since $d\mathcal{L}_q(\mathcal{X}_q^{\text{gC}}) = \|\text{grad } \mathcal{L}_q\|_{\tilde{g}}^2 \geq 0$, we see that $\|\mathcal{X}_q^{\text{gC}}(x)\|_{\tilde{g}}^2 \leq \frac{1}{4} \|\mathcal{X}_q^{\text{gC}}(x)\|_{\tilde{g}}^2$. Therefore, x is a stationary point of $\mathcal{X}_q^{\text{gC}}$. This implies that x is in $B(2R)$. For $\lambda \gg 0$, by the work of Section 7.2, we have that x has L_k^2 distance (and thus L_{k-1}^2 distance) at most $\epsilon/2$ from a stationary point x_λ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$. Since the x_n converge to x in L_{k-1}^2 , they are eventually within L_{k-1}^2 distance ϵ of x_{λ_n} for $n \gg 0$. Since x_{λ_n} is a stationary point of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$, this is a contradiction. \square

However, \mathcal{L}_q does not satisfy (180) in the neighborhoods of stationary points. If it did, then by Lemma 2.21, any stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, i.e., zero of $l + p^\lambda c_q$, would be a critical point of $\mathcal{L}_q|_{W^\lambda}$. We can write the \tilde{g} -gradient of $\mathcal{L}_q|_{W^\lambda}$ as $l + \tilde{p}_g^\lambda c_q$, where \tilde{p}_g^λ is the \tilde{g} -orthogonal projection from W to W^λ . (Compare Remark 3.4.) However, in general, the condition $(l + p^\lambda c_q)(x) = 0$ does not imply $(l + \tilde{p}_g^\lambda c_q)(x) = 0$.

8.2. The function F_λ . To construct the desired function F_λ as in Proposition 8.2, we need to alter \mathcal{L}_q near the stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$.

Let us first introduce some notation. Given a point $x \in W$, its S^1 -orbit can be either a point or a circle. In particular, the stationary points of $\mathcal{X}_q^{\text{gC}}$ come in finitely many such orbits, which we denote by $\mathcal{O}^1, \dots, \mathcal{O}^m$.

Throughout the rest of Section 8, we will fix some $\epsilon > 0$ sufficiently small such that it satisfies the following.

Assumption 8.5.

- (a) *The L_{k-1}^2 distance between any two orbits $\mathcal{O}^j, \mathcal{O}^{j'}$ ($j \neq j'$) is at least 7ϵ ;*
- (b) *If an orbit \mathcal{O}^j consists of irreducibles, then the L^2 norm of a point in \mathcal{O}^j is at least 4ϵ .*

In Section 8.5 we will add another assumption on ϵ . However, Assumption 8.5 above suffices for the results in the current subsection.

With ϵ fixed, we will state our results for λ being sufficiently large. Of course, how large λ is may depend on ϵ .

It follows from Proposition 7.7 and Lemma 7.9 that, for $\lambda \gg 0$, there is a one-to-one correspondence between the orbits of stationary points of $\mathcal{X}_q^{\text{gC}}$, on the one hand, and the orbits of stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ inside $B(2R)$, on the other hand. Further, this correspondence preserves the type of orbits (reducible or irreducible). Let $\mathcal{O}_\lambda^1, \dots, \mathcal{O}_\lambda^m$ be the latter set of orbits, with \mathcal{O}_λ^j corresponding to \mathcal{O}^j . By choosing λ sufficiently large, we can arrange so that, for all j , the orbits \mathcal{O}^j and \mathcal{O}_λ^j are within L_{k-1}^2 distance ϵ of each other. In view of part (a) in Assumption 8.5, this ensures that \mathcal{O}_λ^j and $\mathcal{O}_\lambda^{j'}$ are at least L_{k-1}^2 distance 5ϵ apart, for $j \neq j'$.

Next, consider the neighborhoods $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$ of \mathcal{O}_λ^j , consisting of points at L_{k-1}^2 distance at most 2ϵ from these orbits. Because of our choice of ϵ , all these neighborhoods are disjoint from each other. As an aside, note also that these neighborhoods may well go outside of $B(2R)$, since the latter ball is taken in the L_k^2 metric.

Pick a point x_λ^j on each orbit \mathcal{O}_λ^j . We define functions

$$\omega_\lambda^j : \nu_{2\epsilon}(\mathcal{O}_\lambda^j) \rightarrow S^1$$

as follows. If x_λ^j is reducible, so that $\mathcal{O}_\lambda^j = \{x_\lambda^j\}$, we simply set $\omega_\lambda^j = 1$. If x_λ^j is irreducible, so that \mathcal{O}_λ^j is a circle, we ask that $\omega_\lambda^j(x) \cdot x_\lambda^j$ is the point on \mathcal{O}_λ^j that is at minimal L^2 distance from x . To make sure that ω_λ^j is well-defined, we need to check that this point is unique. The closest point is not unique only for points in the L^2 -orthogonal complement to the plane $\text{Span}(\mathcal{O}_\lambda^j)$. However, if $x \in \nu_{2\epsilon}(\mathcal{O}_\lambda^j)$, then there is some $x' \in \mathcal{O}_\lambda^j$ within L_{k-1}^2 (and hence L^2) distance 2ϵ from x . By part (b) in Assumption 8.5, together with the fact that \mathcal{O}_λ^j and \mathcal{O}^j are within L_{k-1}^2 (and hence L^2) distance ϵ from each other, we see that the L^2 norm of x' is at least 3ϵ . This shows that x' cannot be perpendicular to x , and the claim about the uniqueness of the L^2 -closest point to x follows.

Explicitly, when x_λ^j is irreducible, we can write

$$(181) \quad \omega_\lambda^j(x) = \frac{\text{Re}\langle x, x_\lambda^j \rangle_{L^2} + i \text{Re}\langle x, ix_\lambda^j \rangle_{L^2}}{((\text{Re}\langle x, x_\lambda^j \rangle_{L^2})^2 + (\text{Re}\langle x, ix_\lambda^j \rangle_{L^2})^2)^{1/2}}.$$

Note that the original orbit \mathcal{O}^j is at L_{k-1}^2 distance at most ϵ from \mathcal{O}_λ^j , and hence is contained in $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$. Let x_∞^j be the point in \mathcal{O}^j that is closest in L^2 distance to the chosen basepoint $x_\lambda^j \in \mathcal{O}_\lambda^j$. Then, also x_λ^j is the L^2 -closest point to x_∞^j in \mathcal{O}_λ^j ; in other words, we have

$$\omega_\lambda^j(x_\infty^j) = 1.$$

Let $h : [0, \infty) \rightarrow \mathbb{R}$ be a smooth, non-increasing function such that $h(x) = 1$ for $x \leq 1$ and $h(x) = 0$ for $x \geq 2$. Set

$$H_\lambda^j : W^\lambda \rightarrow [0, 1], \quad H_\lambda^j(x) = h(\epsilon^{-1} d_{L_{k-1}^2}(x, \mathcal{O}_\lambda^j))$$

where $d_{L_{k-1}^2}$ denotes L_{k-1}^2 distance. Note that H_λ^j is identically 0 outside of $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$, and is identically 1 in the smaller neighborhood $\nu_\epsilon(\mathcal{O}_\lambda^j)$.

We now define

$$T_\lambda : W^\lambda \rightarrow W$$

by

$$(182) \quad T_\lambda(x) = x + \sum_{j=1}^m H_\lambda^j(x) \cdot \omega_\lambda^j(x) \cdot (x_\infty^j - x_\lambda^j)$$

and finally set

$$(183) \quad F_\lambda : W^\lambda \rightarrow \mathbb{R}, \quad F_\lambda = \mathcal{L}_q \circ T_\lambda.$$

The function F_λ is the one we will use to prove Proposition 8.2. Before analyzing this function, we give a more qualitative description for the benefit of the reader.

Observe that T_λ and F_λ are S^1 -equivariant, by construction. In the smaller neighborhood $\nu_\epsilon(\mathcal{O}_\lambda^j)$, if we restrict to the affine space perpendicular to \mathcal{O}_λ^j at x_λ^j , we have $\omega_\lambda^j \equiv 1$ there, and hence the map T_λ is given by translation by $x_\infty^j - x_\lambda^j$. In particular, $T_\lambda(x_\lambda^j) = x_\infty^j$. More generally, T_λ takes the orbit \mathcal{O}_λ^j to \mathcal{O}^j . In fact, we can view $\nu_\epsilon(\mathcal{O}_\lambda^j)$ as a disk bundle over \mathcal{O}_λ^j , with the projection map given by taking the L^2 closest point on the orbit. Then, we can say that inside $\nu_\epsilon(\mathcal{O}_\lambda^j)$, the map T_λ consists of fiberwise translations, arranged so that \mathcal{O}_λ^j is taken to \mathcal{O}^j . Further, T_λ is the identity outside $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$, and in the intermediate region $\nu_{2\epsilon}(\mathcal{O}_\lambda^j) \setminus \nu_\epsilon(\mathcal{O}_\lambda^j)$, it is given by some interpolation between fiberwise translation and the identity.

The resulting function $F_\lambda : W^\lambda \rightarrow \mathbb{R}$ agrees with the perturbed CSD functional \mathcal{L}_q outside $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$, whereas near \mathcal{O}_λ^j we have arranged so that the points of \mathcal{O}_λ^j became critical points of F_λ . Effectively, this was accomplished by translating \mathcal{O}_λ^j to \mathcal{O}^j , and using the fact that \mathcal{O}^j consists of stationary points of $\mathcal{X}_q^{\text{gC}}$, i.e., critical points of \mathcal{L}_q .

8.3. Control in the intermediate region. Since F_λ agrees with \mathcal{L}_q at any point with L_{k-1}^2 distance at least 2ϵ from a stationary point of $\mathcal{X}_q^{\text{gC}}$, Lemma 8.4 implies that $dF_\lambda(\mathcal{X}_q^{\text{gC}}) > 0$ in this region. In the current subsection, we will be able to use the same arguments to show that $dF_\lambda(\mathcal{X}_q^{\text{gC}}) > 0$ as long as the distance is at least ϵ , thus gaining control in the intermediate region $\nu_{2\epsilon}(\mathcal{O}_\lambda^j) \setminus \nu_\epsilon(\mathcal{O}_\lambda^j)$. For notation, we will sometimes write $\tilde{g}(x)$ for the inner product given by \tilde{g} on $T_x W$. We begin with a technical lemma.

Lemma 8.6. *Consider a sequence x_n in $B(2R)$ which converges in L_{k-1}^2 to some $x \in W_{k-1}$. Then if $\lambda_n \rightarrow \infty$, $(dF_{\lambda_n})_{x_n}(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}) \rightarrow (d\mathcal{L}_q)_x(\mathcal{X}_q^{\text{gC}})$ in \mathbb{R} .*

Proof. We first compute that for any x' ,

$$\begin{aligned} (dF_\lambda)_{x'}(v) &= (d\mathcal{L}_q)_{T_\lambda(x')}(\mathcal{D}_{x'}T_\lambda)(v) \\ &= \langle \mathcal{X}_q^{\text{gC}}(T_\lambda(x')), \mathcal{D}_{x'}T_\lambda(v) \rangle_{\tilde{g}(x')}. \end{aligned}$$

We claim it suffices to show that $T_{\lambda_n}(x_n) \rightarrow x$ and $(\mathcal{D}_{x_n}T_{\lambda_n})(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}(x_n)) \rightarrow \mathcal{X}_q^{\text{gC}}(x)$, each in L_{k-2}^2 . Indeed, since the $\tilde{g}(x_n)$ -metrics converge to $\tilde{g}(x)$, this will imply that $(dF_{\lambda_n})_{x_n}(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}})$ converges to $\langle \mathcal{X}_q^{\text{gC}}(x), \mathcal{X}_q^{\text{gC}}(x) \rangle_{\tilde{g}(x)}$, which is exactly $(d\mathcal{L}_q)_x(\mathcal{X}_q^{\text{gC}})$.

We begin by analyzing the continuity of T_λ , using (182). Note that $|H_\lambda^j|$ and $|\omega_\lambda^j|$ are bounded above by 1. Since $\lambda_n \rightarrow \infty$, then $x_\infty^j - x_\lambda^j \rightarrow 0$ in L_{k-1}^2 by the discussion after Corollary 7.8. Therefore, we have $T_{\lambda_n}(x_n) \rightarrow x$ in L_{k-1}^2 .

Thus, it remains to analyze $\mathcal{D}_x T_\lambda$. We have that

$$(\mathcal{D}_x T_\lambda)(v) = v + \sum_j (dH_\lambda^j)_x(v) \cdot \omega_\lambda^j(x) \cdot (x_\infty - x_\lambda) + \sum_j H_\lambda^j(x) \cdot \mathcal{D}_x \omega_\lambda^j(v) \cdot (x_\infty - x_\lambda).$$

Note that $|(dH_\lambda^j)_x(v)| \leq (C/\epsilon)\|v\|_{L_{k-1}^2}$ for a constant C independent of x, j , and λ . (More precisely, C is the C^0 -norm of h' .) Also, ω_λ^j is a C^1 -function when restricted to $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$ whose denominator in (181) is bounded below by $3\epsilon^2$, since any point in \mathcal{O}_λ^j has L^2 norm at least 3ϵ . From this, it is easy to obtain bounds

$$(184) \quad |\mathcal{D}_x \omega_\lambda^j(v)| \leq C'\|v\|_{L^2} \leq C'\|v\|_{L_{k-1}^2}$$

independent of $x \in B(2R), j$, and λ . Therefore, $(\mathcal{D}_{x_n} T_{\lambda_n})(v_n) \rightarrow v$ in L_{k-2}^2 for any sequence v_n which converges to v in L_{k-2}^2 and is L_{k-1}^2 bounded. Since $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}(x_n) \rightarrow \mathcal{X}_q^{\text{gC}}(x)$ in L_{k-2}^2 and $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}(x_n)$ is L_{k-1}^2 bounded (because the x_n are L_k^2 bounded), we have that $(\mathcal{D}_{x_n} T_{\lambda_n})(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}(x_n)) \rightarrow \mathcal{X}_q^{\text{gC}}(x)$ in L_{k-2}^2 . This suffices to complete the proof. \square

With the above lemma, we now establish the analogue of Lemma 8.4 for F_λ .

Proposition 8.7. *Fix $\epsilon > 0$ satisfying Assumption 8.5. For $\lambda \gg 0$, we have*

$$(185) \quad \frac{1}{4}\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2 < dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) < 4\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2$$

at any point in $W^\lambda \cap B(2R)$ which is at L_{k-1}^2 distance at least ϵ from any stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}}$.

Proof. Suppose that the conclusion is not true. Then, there exists a sequence $\lambda_n \rightarrow \infty$ and a sequence $x_n \in W^{\lambda_n} \cap B(2R)$ which are L_{k-1}^2 distance at least ϵ from any stationary point of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ and $(dF_{\lambda_n})_{x_n}(\mathcal{X}_{q^{\lambda_n}}^{\text{gC}})$ violate (185). Then, there exists a subsequence of the x_n which converges in L_{k-1}^2 to an element $x \in W_{k-1}$. By Lemma 8.6, we see that $(d\mathcal{L}_q)_x(\mathcal{X}_q^{\text{gC}})$ must be at most $\frac{1}{4}\|\mathcal{X}_q^{\text{gC}}(x)\|_{\tilde{g}}^2$ or at least $4\|\mathcal{X}_q^{\text{gC}}(x)\|_{\tilde{g}}^2$. Since $(d\mathcal{L}_q)(\mathcal{X}_q^{\text{gC}}) = \|\text{grad } \mathcal{L}_q\|^2$, we get that x is a stationary point of $\mathcal{X}_q^{\text{gC}}$, and thus $x \in B(2R) \subset W_k$. Since the x_n converge to x in L_{k-1}^2 , they are eventually within L_{k-1}^2 distance $\epsilon/2$ of a stationary point of $\mathcal{X}_q^{\text{gC}}$. For $\lambda \gg 0$, there is a stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ within L_{k-1}^2 distance $\epsilon/2$ of x . This contradicts the fact that the x_n are not within L_{k-1}^2 distance ϵ of a stationary point of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$. \square

8.4. The L^2 and \tilde{g} metrics. A more detailed analysis will be needed to prove the inequality (179) in neighborhoods of the stationary points. This will be done in Section 8.5. As a preliminary step, since $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is an L^2 approximation to the \tilde{g} -gradient of $\mathcal{L}_q|_{W^\lambda}$, we will prove a few results relating the L^2 and \tilde{g} metrics.

Let us recall the definition of \tilde{g} from Section 3.3. For $x = (a, \phi) \in W$ and $(b, \psi), (b', \psi') \in T_{(a, \phi)}W$, we have

$$\langle (b, \psi), (b', \psi') \rangle_{\tilde{g}} = \langle \Pi_{(a, \phi)}^{\text{elC}}(b, \psi), \Pi_{(a, \phi)}^{\text{elC}}(b', \psi') \rangle_{L^2},$$

where Π^{elC} is the enlarged local Coulomb projection. Let us also recall the formula for this projection:

$$(186) \quad \Pi_{(a, \phi)}^{\text{elC}}(b, \psi) := (b - d\zeta, \psi + \zeta\phi),$$

where $\zeta : Y \rightarrow i\mathbb{R}$ is determined (for $\phi \neq 0$) by the conditions $\int_Y \zeta = 0$ and

$$(187) \quad \Delta\zeta + |\phi|^2\zeta - \mu_Y(|\phi|^2\zeta) = -i\operatorname{Re}\langle i\phi, \psi \rangle + i\mu_Y(\operatorname{Re}\langle i\phi, \psi \rangle).$$

The last equality is Equation (31), where we used the fact that $d^*b = 0$ for $(b, \psi) \in T_x W$.

Lemma 8.8. *There is a constant $K > 0$ such that, for all $x = (a, \phi) \in B(2R) \subset W_k$ and $(b, \psi) \in \mathcal{T}_{0,x}^{\text{gC}} \cong W_0$, if ζ is the function in (186), then*

$$(188) \quad \|\zeta\|_{L_1^2} \leq K \cdot \|\psi\|_{L_{-1}^2}.$$

Proof. This follows from Lemma 3.2 by considering the operator E_ϕ from L_1^2 to L_{-1}^2 . Note that K can be taken to be a constant independent of ϕ , because the L_k^2 norm of ϕ is bounded. \square

Since $\|\psi\|_{L_{-1}^2} \leq \|\psi\|_{L^2}$, Lemma 8.8 also gives L_1^2 bounds on ζ in terms of L^2 bounds on ψ .

Our next result is about the equivalence between the L^2 and \tilde{g} metrics.

Proposition 8.9. *There is a constant $C_0 > 0$ such that, for all $x = (a, \phi) \in B(2R) \subset W_k$ and $(b, \psi) \in \mathcal{T}_{0,x}^{\text{gC}} \cong W_0$, we have*

$$(189) \quad C_0^{-1} \cdot \|(b, \psi)\|_{L^2} \leq \|(b, \psi)\|_{\tilde{g}(x)} \leq C_0 \cdot \|(b, \psi)\|_{L^2}.$$

Proof. We begin with the second inequality in (189). Then, with ζ as in (186), we have

$$\|(b, \psi)\|_{\tilde{g}(x)} \leq \|(b, \psi)\|_{L^2} + \|(d\zeta, \zeta\phi)\|_{L^2}.$$

We claim that the right hand side is bounded by a constant times $\|(b, \psi)\|_{L^2}$. This follows by the L_1^2 control on ζ from Lemma 8.8 and the fact that the condition $x \in B(2R)$ gives L_k^2 (and hence C^0) bounds on ϕ .

To prove the first inequality, note that if we view Π_x^{elC} as an isomorphism from the global Coulomb slice to the extended local Coulomb slice, then its inverse is the infinitesimal global Coulomb projection, $(\Pi_*^{\text{gC}})_x$, from (24). Thus, if we switch notation and now let (b, ψ) be a vector in the extended local Coulomb slice, the first inequality in (189) can be re-written as

$$\|(\Pi_*^{\text{gC}})_x(b, \psi)\|_{L^2} \leq C_0 \cdot \|(b, \psi)\|_{L^2}.$$

This holds by Lemma 5.12. \square

Our next goal is to compare the L^2 - to \tilde{g} -orthogonal projections from W to W^λ . Since λ is of the form λ_i^\bullet , we have that p^λ is the L^2 -orthogonal projection to W^λ , and we write $p_{\tilde{g}(x)}^\lambda$ for the $\tilde{g}(x)$ -orthogonal projection to W^λ . (Compare Remark 3.4.)

Proposition 8.10. *There is a constant $C_1 > 0$ such that, for all $x \in B(2R) \subset W_k$, we have*

$$\|p_{\tilde{g}(x)}^\lambda - p^\lambda\| \leq \frac{C_1}{\lambda},$$

where $p_{\tilde{g}(x)}^\lambda$ and p^λ are viewed as operators from the L^2 completion W_0 of W to itself.

We first need a refinement of Lemma 8.8. Let us denote by $(W^\lambda)_0^\perp$ the L^2 -orthogonal complement to W^λ inside W_0 . In other words, $(W^\lambda)_0^\perp$ is the L^2 span of the eigenvectors of l with eigenvalues at least λ in absolute value.

Lemma 8.11. *There is a constant $C_2 > 0$ such that, for all $x = (a, \phi) \in B(2R) \subset W_k$ and $(b, \psi) \in (W^\lambda)_0^\perp$, if ζ is the function in (186), then*

$$\|\zeta\|_{L^2_1} \leq \frac{C_2}{\lambda} \|\psi\|_{L^2}.$$

Proof. Write $\psi = \sum \psi_\kappa$, where ψ_κ are eigenvectors of the Dirac operator D , with eigenvalues κ such that $|\kappa| > \lambda$. Set

$$(190) \quad D^{-1}(\psi) := \sum \kappa^{-1} \psi_\kappa.$$

Note that $D(D^{-1}(\psi)) = \psi$. Since D is continuous from L^2 spinors to L^2_{-1} spinors, we obtain

$$\|\psi\|_{L^2_{-1}} \leq K' \|D^{-1}\psi\|_{L^2} \leq \frac{K'}{\lambda} \|\psi\|_{L^2},$$

for some constant K' . Together with Lemma 8.8, this gives the desired inequality. \square

Proof of Proposition 8.10. We need to check that, for all $v \in W_0$,

$$\|p_{\tilde{g}(x)}^\lambda(v) - p^\lambda(v)\|_{L^2} \leq \frac{C_1}{\lambda} \|v\|_{L^2},$$

where C_1 is independent of x, v , and λ .

First, if $v \in W^\lambda$, then note that $p_{\tilde{g}(x)}^\lambda(v) = p^\lambda(v)$ and the claim is trivial. Therefore, we can assume that $v \in (W^\lambda)_0^\perp$. Thus, we would like to bound

$$\|p_{\tilde{g}(x)}^\lambda(v) - p^\lambda(v)\|_{L^2} = \|p_{\tilde{g}(x)}^\lambda(v)\|_{L^2}.$$

We write $x = (a, \phi)$ and

$$w = p_{\tilde{g}(x)}^\lambda(v) \in W^\lambda.$$

Recall from Proposition 8.9 that there exists a constant C_0 , independent of x , such that $\|v\|_{L^2} \leq C_0 \|v\|_{\tilde{g}(x)}$ and similarly for w . In this case, we have that

$$\|p_{\tilde{g}(x)}^\lambda(v)\|_{L^2}^2 \leq C_0^2 \|p_{\tilde{g}(x)}^\lambda(v)\|_{\tilde{g}(x)}^2 = C_0^2 \langle v, p_{\tilde{g}(x)}^\lambda(v) \rangle_{\tilde{g}(x)}.$$

Therefore, to obtain the desired bounds in the proposition, it suffices to prove that there exists a constant $K_0 > 0$, independent of $x \in B(2R)$ and $v \in (W^\lambda)_0^\perp$, such that

$$(191) \quad |\langle v, w \rangle_{\tilde{g}(x)}| \leq \frac{K_0}{\lambda} \|v\|_{L^2} \|w\|_{L^2}.$$

We now focus on proving this inequality.

For notation, let $v = (b, \psi)$ and $w = (b', \psi')$. We write $\Pi_x^{\text{elC}}(v) = (b - d\zeta, \psi + \zeta\phi)$ and $\Pi_x^{\text{elC}}(w) = (b' - d\zeta', \psi' + \zeta'\phi)$. Since $w = p_{\tilde{g}(x)}^\lambda(v)$, we have that v and w are L^2 orthogonal. From this, we obtain

$$(192) \quad \begin{aligned} \langle v, w \rangle_{\tilde{g}(x)} &= \langle \Pi_x^{\text{elC}}(v), \Pi_x^{\text{elC}}(w) \rangle_{L^2} \\ &= -\langle d\zeta, b' \rangle_{L^2} - \langle b, d\zeta' \rangle_{L^2} + \langle d\zeta, d\zeta' \rangle_{L^2} + \langle \zeta\phi, \psi' \rangle_{L^2} \\ &\quad + \langle \psi, \zeta'\phi \rangle_{L^2} + \langle \zeta\phi, \zeta'\phi \rangle_{L^2} \end{aligned}$$

$$(193) \quad = \langle d\zeta, d\zeta' \rangle_{L^2} + \langle \zeta\phi, \psi' \rangle_{L^2} + \langle \psi, \zeta'\phi \rangle_{L^2} + \langle \zeta\phi, \zeta'\phi \rangle_{L^2},$$

where the third equality comes from the fact that $d^*b = d^*b' = 0$. Let us now collect the relevant bounds on ζ and ζ' . By Lemma 8.11, we have that $\|\zeta\|_{L^2}$ and $\|d\zeta\|_{L^2}$ are bounded above by $\frac{C_2}{\lambda} \|\psi\|_{L^2}$ for some constant $C_2 > 0$. We do not have L^2 bounds on ζ' or $d\zeta'$ in

terms of $\frac{1}{\lambda}$, since w is not an element of $(W^\lambda)_0^\perp$. However, Lemma 8.8 still guarantees $\|\zeta'\|_{L^2}$ and $\|d\zeta'\|_{L^2}$ are bounded above by $K \cdot \|\psi'\|_{L^2}$, independent of λ and $x \in B(2R)$. This gives:

$$(194) \quad |\langle d\zeta, d\zeta' \rangle_{L^2}| \leq \frac{C_2 \cdot K}{\lambda} \|\psi\|_{L^2} \|\psi'\|_{L^2} \leq \frac{C_2 \cdot K}{\lambda} \|v\|_{L^2} \|w\|_{L^2}.$$

To establish (191), it remains to bound the other three terms in (193).

Since $x \in B(2R) \subset W_k$, we have a uniform bound on $\|\phi\|_{C^0}$ independent of x , denoted K_2 . Therefore, using a similar argument as for (194), we obtain bounds

$$\begin{aligned} |\langle \zeta\phi, \psi' \rangle_{L^2}| &\leq \frac{C_2 \cdot K_2}{\lambda} \|v\|_{L^2} \|w\|_{L^2}, \\ |\langle \zeta\phi, \zeta'\phi \rangle_{L^2}| &\leq \frac{C_2 \cdot K \cdot K_2}{\lambda} \|v\|_{L^2} \|w\|_{L^2}. \end{aligned}$$

Thus, the proof will be complete if we obtain similar bounds for $\langle \psi, \zeta'\phi \rangle_{L^2}$. We must be careful here, since we do not have bounds on ζ' in terms of $\frac{1}{\lambda}$. To handle this, we write

$$\begin{aligned} \langle \psi, \zeta'\phi \rangle_{L^2} &= \langle D(D^{-1}\psi), \zeta'\phi \rangle_{L^2} \\ &= \langle D^{-1}\psi, D(\zeta'\phi) \rangle_{L^2}, \end{aligned}$$

where D^{-1} is defined as in (190). Since $(b, \psi) \in (W^\lambda)_0^\perp$, we have that $\|D^{-1}\psi\|_{L^2} \leq \frac{1}{\lambda} \|\psi\|_{L^2}$. Finally, since ζ' is L_1^2 -bounded in terms of $\|\psi\|_{L^2}$, we obtain L_1^2 -bounds on $\zeta'\phi$ (independent of x) by Sobolev multiplication. In other words, there exists a constant K_3 such that $\|D(\zeta'\phi)\|_{L^2} \leq K_3 \|\psi'\|_{L^2}$, since D is continuous from L_1^2 to L^2 . Thus, we conclude that

$$|\langle \psi, \zeta'\phi \rangle_{L^2}| \leq \frac{K_3}{\lambda} \|v\|_{L^2} \|w\|_{L^2},$$

which completes the proof. \square

8.5. Control near the stationary points. We are now ready to prove (179) in the neighborhoods $\nu_\epsilon(\mathcal{O}_\lambda^j)$ of the orbits of stationary points.

Consider the stationary point x_∞^j of $\mathcal{X}_q^{\text{gC}}$. Since x_∞^j is non-degenerate, we have that when restricted to the anticircular global Coulomb slice, $l + \mathcal{D}_{x_\infty^j} c_q$ is an invertible, self-adjoint operator with respect to the metric \tilde{g} . We will use this to show that $(l + p^\lambda \mathcal{D}_{x_\lambda^j} c_q)(y)$ grows at least linearly in y when y is L^2 orthogonal to the S^1 -orbit of $x_\lambda^j = (a_\lambda^j, \phi_\lambda^j)$; this will be key for establishing the inequalities (179) near x_λ^j .

Lemma 8.12. *There exists a constant $\mu_j > 0$ with the following property. For $\lambda \gg 0$ and any $y \in W_1$ with $\text{Re}\langle y, (0, i\phi_\lambda^j) \rangle_{L^2} = 0$, we have*

$$(195) \quad \|(l + p^\lambda \mathcal{D}_{x_\lambda^j} c_q)(y)\|_{L^2} \geq \mu_j \|y\|_{L_1^2}.$$

Of course, this lemma implies that we have bounds $\|(l + p^\lambda \mathcal{D}_{x_\lambda^j} c_q)(y)\|_{L^2} \geq \mu_j \|y\|_{L^2}$ as well.

Proof. For notational convenience, we omit the index j from the argument. We suppose that this is not true. Then, we can find a sequence $y_n \in W^{\lambda_n}$ with $\|y_n\|_{L_1^2} = 1$ and $\text{Re}\langle y_n, (0, i\phi_{\lambda_n}) \rangle_{L^2} = 0$ and a sequence $\lambda_n \rightarrow \infty$ such that $\|(l + p^{\lambda_n} \mathcal{D}_{x_{\lambda_n}} c_q)(y_n)\|_{L^2} \rightarrow 0$.

Extract a subsequence for which y_n converges in L^2 to some $y \in W_0$. Since $x_{\lambda_n} \rightarrow x_\infty$ in L_k^2 , the continuity of $\mathcal{D}c_q$ (Lemma 6.11) implies that

$$p^{\lambda_n} \mathcal{D}_{x_{\lambda_n}} c_q(y_n) \rightarrow \mathcal{D}_{x_\infty} c_q(y) \text{ in } L^2.$$

We see that $l(y_n)$ converges in L^2 . Because the L_{-1}^2 limit of $l(y_n)$ is $l(y)$, we in fact have that $l(y_n)$ converges to $l(y)$ in L^2 and thus $y_n \rightarrow y$ in L_1^2 , and $y \neq 0$. However, since $(l + p^{\lambda_n} \mathcal{D}_{x_{\lambda_n}} c_q)(y_n)$ converges to 0, we see that $(l + \mathcal{D}_{x_\infty} c_q)(y) = 0$. Bootstrapping further shows that y is actually an element of W_k .

First, suppose that x_∞ is reducible. In this case, we have contradicted the non-degeneracy of x_∞ , as $l + \mathcal{D}_{x_\infty} c_q : W_k \rightarrow W_{k-1}$ is invertible (where we are using non-degeneracy in the blow-down). Now, suppose that x_∞ is an irreducible stationary point of $\mathcal{X}_q^{\text{gC}}$. In this case, we can only say that $l + \mathcal{D}_{x_\infty} c_q$ is injective on the $\tilde{g}(x_\infty)$ -orthogonal complement of the S^1 orbit of x_∞ . Thus, we have that y must have some component tangent to the S^1 orbit of x_λ , i.e., a multiple of $(0, i\phi_\infty)$. Because $\text{Re}\langle y_n, (0, i\phi_{\lambda_n}) \rangle_{L^2} = 0$ for all n , we see that $\text{Re}\langle y, (0, i\phi_\infty) \rangle_{L^2} = 0$ as well. This is a contradiction. \square

Recall that $\epsilon > 0$ was chosen in Section 8.2 to be sufficiently small, depending on two requirements from Assumption 8.5. We need an additional requirement. For each $j = 1, \dots, m$, let μ_j be the constant as in Lemma 8.12.

Assumption 8.13. *For each $j = 1, \dots, m$ and $x \in \nu_{3\epsilon}(\mathcal{O}^j)$, we have*

$$\|\mathcal{D}_x c_q - \mathcal{D}_{x_\infty^j} c_q\| \leq \frac{\mu_j}{40C_0^2}.$$

Here, the operator norm is taken in $\text{Hom}(W_0, W_0)$, and C_0 is the constant from Proposition 8.9.

Note that the existence of such an ϵ is guaranteed by the continuity of $\mathcal{D}c_q$, as Lemma 6.11, applied for Sobolev index $k-1$ instead of k , implies that

$$\mathcal{D}c_q : W_{k-1} \rightarrow \text{Hom}(W_0, W_0)$$

is continuous.

Proposition 8.14. *Fix $\epsilon > 0$ satisfying Assumptions 8.5 and 8.13. Then, for all $\lambda \gg 0$, we have*

$$(196) \quad \frac{1}{4} \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2 \leq dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \leq 4 \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2$$

at any point $x \in W^\lambda \cap B(2R)$ which is at L_{k-1}^2 distance at most ϵ from a stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}}$.

Proof. In Section 8.2 we observed that, because of Assumption 8.5 (a) on ϵ , the S^1 -orbits \mathcal{O}_λ^j of stationary points x_λ^j are at least L_{k-1}^2 distance 5ϵ apart from each other. Hence, their neighborhoods $\nu_{2\epsilon}(\mathcal{O}_\lambda^j)$ are disjoint.

Let x be in $\nu_\epsilon(\mathcal{O}_\lambda^j)$ for some j . Then, the only contribution to the summation in Equation (182) is from that j and, moreover, we have $H_\lambda^j(x) = 1$. For simplicity, we will omit the index j from x_λ^j , x_∞^j , \mathcal{O}^j , \mathcal{O}_λ^j , and ω_λ^j for the rest of this section. Thus,

$$T_\lambda(x) = x + \omega_\lambda(x) \cdot (x_\infty - x_\lambda).$$

We now expand $dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}})$:

$$\begin{aligned} (dF_\lambda)_x(\mathcal{X}_{q^\lambda}^{\text{gC}}(x)) &= (d\mathcal{L}_q)_{T_\lambda(x)}(\mathcal{D}_x T_\lambda)(\mathcal{X}_{q^\lambda}^{\text{gC}}(x)) \\ &= \langle (l + c_q)(T_\lambda(x)), (\mathcal{D}_x T_\lambda)(l + p^\lambda c_q)(x) \rangle_{\tilde{g}(T_\lambda(x))} \\ &= \langle (l + c_q)(T_\lambda(x)), (l + p^\lambda c_q)(x) + (\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q)(x) \cdot (x_\infty - x_\lambda) \rangle_{\tilde{g}(T_\lambda(x))}. \end{aligned}$$

By the S^1 -equivariance of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ and the S^1 -invariance of F_λ , it suffices to show that (196) holds when $\omega_\lambda(x) = 1$, that is, when the L^2 -closest point to x on \mathcal{O}_λ is exactly x_λ . (Of course, we automatically have $\omega_\lambda(x) = 1$ if x_λ is reducible.) Further, note that if x_λ is reducible, then $\mathcal{D}_x \omega_\lambda = 0$.

We will analyze $dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}})$ by linearizing some of the terms in the inner product above about x_λ . We begin by linearizing $(l + p^\lambda c_q)(x)$. Let $y = x - x_\lambda$. Since $\omega_\lambda(x) = 1$, we have that y is orthogonal to the S^1 orbit of x_λ . Since x_λ is a stationary point, we have $(l + p^\lambda c_q)(x_\lambda) = 0$. Therefore, the difference between $(l + p^\lambda c_q)(y)$ and $(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(x)$ is given by

$$\begin{aligned} (197) \quad (l + p^\lambda c_q)(x_\lambda + y) - (l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y) &= l(x_\lambda) + p^\lambda(c_q(x_\lambda + y)) - p^\lambda(\mathcal{D}_{x_\lambda} c_q(y)) \\ &= -p^\lambda(c_q(x_\lambda) + c_q(x_\lambda + y) - \mathcal{D}_{x_\lambda} c_q(y)) \end{aligned}$$

$$(198) \quad = \int_0^1 p^\lambda(\mathcal{D}_{x_\lambda + ty} c_q - \mathcal{D}_{x_\lambda} c_q)(y) dt.$$

We denote the term in (198) by $\tilde{S}_\lambda(y)$.

Similarly, we would like to linearize $(l + c_q)(T_\lambda(x))$. We first recall that $T_\lambda(x_\lambda) = x_\infty$, which is a stationary point of $l + c_q$. Second, since $\omega_\lambda(x) = 1$, we have that $x - x_\lambda$ is real L^2 orthogonal to ix_λ , and hence $(\mathcal{D}_{x_\lambda + ty} T_\lambda)(y) = y$ for any $t \in [0, 1]$. Therefore, we can write

$$\begin{aligned} (l + c_q)(T_\lambda(x_\lambda + y)) &= (l + c_q)(T_\lambda(x_\lambda + y)) - (l + c_q)(T_\lambda(x_\lambda)) \\ &= \int_0^1 (l + \mathcal{D}_{T_\lambda(x_\lambda + ty)} c_q) \circ (\mathcal{D}_{x_\lambda + ty} T_\lambda)(y) dt \\ &= \int_0^1 (l + \mathcal{D}_{T_\lambda(x_\lambda + ty)} c_q)(y) dt \\ &= \int_0^1 (l + \mathcal{D}_{x_\infty + ty} c_q)(y) dt. \end{aligned}$$

We obtain

$$(199) \quad (l + c_q)(T_\lambda(x_\lambda + y)) - (l + \mathcal{D}_{x_\infty} c_q)(y) = \int_0^1 (\mathcal{D}_{x_\infty + ty} c_q - \mathcal{D}_{x_\infty} c_q)(y) dt.$$

We let $\tilde{R}_\lambda(y)$ denote the term on the right hand side of (199).

Using these two linearizations and the fact that $\langle (1 - p_{\tilde{g}(x')}^\lambda)(u), v \rangle_{\tilde{g}(x')} = 0$ for any $v \in W^\lambda$ and $u, x' \in W_0$, we obtain

$$\begin{aligned} (200) \quad (dF_\lambda)_x(\mathcal{X}_{q^\lambda}^{\text{gC}}(x)) &= \langle (l + p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\infty} c_q)(y) + \tilde{R}_\lambda(y), \\ &\quad (l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y) + \tilde{S}_\lambda(y) + (\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q)(x) \cdot (x_\infty - x_\lambda) \rangle_{\tilde{g}(T_\lambda(x))}, \end{aligned}$$

where $\tilde{R}_\lambda(y)$ and $\tilde{S}_\lambda(y)$ are as defined above. At this point, it is still too difficult to compare the linear terms to understand this inner product. Therefore, we will alter our expressions so that the leading terms align. Write

$$(201) \quad (dF_\lambda)_x(\mathcal{X}_{q^\lambda}^{\text{gC}}(x)) = \langle (l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y) + R_\lambda(y), (l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y) + S_\lambda(y) \rangle_{\tilde{g}(T_\lambda(x))},$$

where

$$(202) \quad R_\lambda(y) = \tilde{R}_\lambda(y) + (p^\lambda_{\tilde{g}(T_\lambda(x))} \mathcal{D}_{x_\infty} c_q - p^\lambda \mathcal{D}_{x_\lambda} c_q)(y),$$

$$(203) \quad S_\lambda(y) = \tilde{S}_\lambda(y) + (\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q)(x) \cdot (x_\infty - x_\lambda).$$

In Lemmas 8.15, 8.16, and 8.17 below, we will prove that for our choice of ϵ as in Assumption 8.13 and for $\lambda \gg 0$,

$$(204) \quad \|\tilde{R}_\lambda(y)\|_{L^2} \leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}$$

$$(205) \quad \|\tilde{S}_\lambda(y)\|_{L^2} \leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}$$

$$(206) \quad \|R_\lambda(y)\|_{L^2} \leq \frac{1}{10C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}$$

$$(207) \quad \|S_\lambda(y)\|_{L^2} \leq \frac{1}{10C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}$$

where C_0 is the constant of equivalency between the \tilde{g} - and L^2 -metrics from Proposition 8.9. This will imply the analogous inequalities with L^2 replaced by $\tilde{g}(T_\lambda(x))$ and without the C_0^2 term. Let us see why this will complete the proof of the proposition.

Note that in an inner product space with vectors u, v, w, \tilde{w} such that

$$(208) \quad \|v\| \leq \frac{1}{10}\|u\|, \quad \|w\| \leq \frac{1}{10}\|u\|, \quad \|\tilde{w}\| \leq \frac{1}{20}\|u\|,$$

we have the inequalities

$$\begin{aligned} \frac{19}{20}\|u\| &\leq \|u + \tilde{w}\| \leq \frac{21}{20}\|u\| \\ \left(\frac{9}{10}\right)^2 \|u\|^2 &\leq \langle u + v, u + w \rangle \leq \left(\frac{11}{10}\right)^2 \|u\|^2, \end{aligned}$$

from which we can deduce

$$\frac{1}{2}\|u + \tilde{w}\|^2 \leq \langle u + v, u + w \rangle \leq 2\|u + \tilde{w}\|^2.$$

For our case, we take

$$u = (l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y), \quad v = R_\lambda(y), \quad w = S_\lambda(y), \quad \tilde{w} = \tilde{S}_\lambda(y),$$

and $\tilde{g}(T_\lambda(x))$ as the inner product. Note that $u + \tilde{w} = \mathcal{X}_{q^\lambda}^{\text{gC}}(x)$. The inequalities (205)-(207) imply that (208) holds and thus

$$\frac{1}{2}\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}(T_\lambda(x))}^2 \leq dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \leq 2\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}(T_\lambda(x))}^2.$$

Since $T_\lambda(x) \rightarrow x$, the constant of equivalency between the $\tilde{g}(T_\lambda(x))$ - and $\tilde{g}(x)$ -metrics is at most $\sqrt{2}$ for $\lambda \gg 0$. From this, we obtain that

$$\frac{1}{4}\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2 \leq dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \leq 4\|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{\tilde{g}}^2,$$

completing the proof. \square

The rest of the section is now devoted to proving (204)- (207). We do this by a series of lemmas.

Lemma 8.15. *For $\lambda \gg 0$ and any $y \in W_1$ orthogonal to the S^1 orbit of x_λ , we have the following inequalities*

$$\begin{aligned}\|\tilde{R}_\lambda(y)\|_{L^2} &\leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2} \\ \|\tilde{S}_\lambda(y)\|_{L^2} &\leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}.\end{aligned}$$

Proof. We will prove the desired inequality for the case of \tilde{S}_λ . The case of \tilde{R}_λ is similar. By Lemma 8.12, $\|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2} \geq \mu \|y\|_{L^2}$. Therefore, it suffices to establish the upper bound

$$(209) \quad \|\tilde{S}_\lambda(y)\|_{L^2} \leq \frac{\mu}{20C_0^2} \|y\|_{L^2}.$$

Since $\|y\|_{L_{k-1}^2} \leq \epsilon$ and $\|x_\lambda - x_\infty\|_{L_{k-1}^2} \leq \epsilon$, we get that $x_\lambda + ty$ is in $\nu_{3\epsilon}(\mathcal{O})$ for all $t \in [0, 1]$. Assumption 8.13 implies that $\|\mathcal{D}_{x_\lambda + ty} c_q - \mathcal{D}_{x_\lambda} c_q\| \leq \frac{\mu}{20C_0^2}$. We obtain

$$\begin{aligned}\left\| \int_0^1 p^\lambda (\mathcal{D}_{x_\lambda + ty} c_q - \mathcal{D}_{x_\lambda} c_q)(y) dt \right\|_{L^2} &\leq \left(\int_0^1 \|p^\lambda\| \cdot \|\mathcal{D}_{x_\lambda + ty} c_q - \mathcal{D}_{x_\lambda} c_q\| dt \right) \cdot \|y\|_{L^2} \\ &\leq \frac{\mu}{20C_0^2} \|y\|_{L^2},\end{aligned}$$

which establishes (209). \square

Lemma 8.16. *For $\lambda \gg 0$, the inequality (206) holds.*

Proof. Using Lemma 8.15 and the definition of R_λ in (202), it suffices to prove that for $\lambda \gg 0$,

$$\|(p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\infty} c_q)(y) - (p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2} \leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}.$$

The proof will be similar to that of Lemma 8.15. By Lemma 8.12, $\|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2} \geq \mu \|y\|_{L^2}$. Therefore, it suffices to establish the upper bound

$$(210) \quad \|p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\infty} c_q - p^\lambda \mathcal{D}_{x_\lambda} c_q\| \leq \frac{\mu}{20C_0^2},$$

where the norm is computed as an operator from W_0 to W_0 . In view of Proposition 8.9, we have that as an operator from L^2 to L^2 , $\|p_{\tilde{g}(T_\lambda(x))}^\lambda\| \leq C_0^2$, as $p_{\tilde{g}}^\lambda$ is a \tilde{g} -orthogonal projection. Therefore,

$$\begin{aligned}\|p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\infty} c_q - p^\lambda \mathcal{D}_{x_\lambda} c_q\| &\leq \|p_{\tilde{g}(T_\lambda(x))}^\lambda (\mathcal{D}_{x_\infty} c_q - \mathcal{D}_{x_\lambda} c_q)\| + \|p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\lambda} c_q - p^\lambda \mathcal{D}_{x_\lambda} c_q\| \\ &\leq C_0^2 \|\mathcal{D}_{x_\infty} c_q - \mathcal{D}_{x_\lambda} c_q\| + \|p_{\tilde{g}(T_\lambda(x))}^\lambda - p^\lambda\| \cdot \|\mathcal{D}_{x_\lambda} c_q\|.\end{aligned}$$

Since $x_\lambda \rightarrow x_\infty$ in L_k^2 norm, the continuity of $\mathcal{D}c_q$ implies that we have uniform bounds on $\|\mathcal{D}_{x_\lambda} c_q\|$ independent of λ ; further, this implies that $\|\mathcal{D}_{x_\infty} c_q - \mathcal{D}_{x_\lambda} c_q\| \rightarrow 0$ as $\lambda \rightarrow \infty$. Finally, by Proposition 8.10, $\|p_{\tilde{g}(T_\lambda(x))}^\lambda - p^\lambda\| \rightarrow 0$ as $\lambda \rightarrow \infty$. Thus, we conclude that for $\lambda \gg 0$, the operator norm of $p_{\tilde{g}(T_\lambda(x))}^\lambda \mathcal{D}_{x_\infty} c_q - p^\lambda \mathcal{D}_{x_\lambda} c_q$ is at most $\frac{\mu}{40C_0^2}$. This proves (210), which we saw earlier was sufficient to complete the proof. \square

Lemma 8.17. *For $\lambda \gg 0$, the inequality (207) holds.*

Proof. Using Lemma 8.15 and the definition of S_λ in (203), it suffices to prove that for $\lambda \gg 0$,

$$(211) \quad \|(\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q(x)) \cdot (x_\infty - x_\lambda)\|_{L^2} \leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}.$$

By the definition of \tilde{S}_λ , we see that

$$(\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q(x)) = (\mathcal{D}_x \omega_\lambda)((l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y) + \tilde{S}_\lambda(y)).$$

Recall from (184), $|(\mathcal{D}_x \omega_\lambda)(v)| \leq C' \|v\|_{L^2}$, for a constant C' independent of $x \in B(2R)$. Combining this inequality with Lemma 8.15, we have

$$|(\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q(x))| \leq C' \left(1 + \frac{1}{20C_0^2}\right) \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}.$$

Since $\|x_\infty - x_\lambda\|_{L^2} \rightarrow 0$, we can choose $\lambda \gg 0$ such that

$$|(\mathcal{D}_x \omega_\lambda)(l + p^\lambda c_q(x))| \cdot \|x_\infty - x_\lambda\|_{L^2} \leq \frac{1}{20C_0^2} \|(l + p^\lambda \mathcal{D}_{x_\lambda} c_q)(y)\|_{L^2}.$$

This establishes (211), and the proof is complete. \square

8.6. Proposition 8.2 and its consequences. Proposition 8.2 now follows by combining Propositions 8.7 and 8.14. Let us also give a few corollaries, which will prove useful in the next sections.

From Proposition 8.2, together with the equivalence between the \tilde{g} and L^2 metrics (Proposition 8.9), we obtain

Corollary 8.18. *There exists $C_0 > 0$ such that for $\lambda \gg 0$,*

$$(212) \quad \frac{1}{4C_0} \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{L^2}^2 \leq dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \leq 4C_0 \|\mathcal{X}_{q^\lambda}^{\text{gC}}\|_{L^2}^2$$

at any $x \in W^\lambda \cap B(2R)$.

We have shown that $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient vector field as in Definition 2.18 (on a non-compact manifold, as discussed in Section 2.8). Lemma 2.22 implies the following:

Corollary 8.19. *Let $\gamma : \mathbb{R} \rightarrow W^\lambda \cap B(2R)$ be a flow line of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, for $\lambda \gg 0$. Then, $\lim_{t \rightarrow -\infty} [\gamma(t)]$ and $\lim_{t \rightarrow +\infty} [\gamma(t)]$ exist in $(W^\lambda \cap B(2R))/S^1$, and they are both projections of stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$.*

There is also the similar result in the blow-up, which is a consequence of Lemma 2.23:

Corollary 8.20. *Let $[\gamma] : \mathbb{R} \rightarrow (W^\lambda \cap B(2R))^\sigma / S^1$ be a flow line of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$, for $\lambda \gg 0$. Then, $\lim_{t \rightarrow -\infty} [\gamma(t)]$ and $\lim_{t \rightarrow +\infty} [\gamma(t)]$ exist in $(W^\lambda \cap B(2R))^\sigma / S^1$, and they are both projections of stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$.*

In fact, one can say more about the limiting behavior of trajectories of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$. Recall that in classical Morse theory, trajectories converge to stationary points with exponential decay (see for example [AD14, Section 10.2.b]). A similar argument works for Morse equivariant quasi-gradients as well. In particular, one obtains that trajectories $[\gamma] : \mathbb{R} \rightarrow (W^\lambda \cap$

$B(2R)^\sigma)/S^1$ of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ from $[x_\lambda]$ to $[y_\lambda]$ must be contained in $\mathcal{B}_k^{\text{gC},\tau}([x_\lambda], [y_\lambda])$. In Section 12 we will compute this exponential decay in terms of F_λ more explicitly and obtain bounds uniform in λ .

If $I \subset \mathbb{R}$ is an interval and $\gamma : I \rightarrow W^\lambda \cap B(2R)$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, its energy is defined¹³ to be

$$(213) \quad \mathcal{E}(\gamma) = \int_I \left\| \frac{d\gamma}{dt} \right\|_{L^2(Y)}^2 dt = \int_I \|\mathcal{X}_{q^\lambda}^{\text{gC}}(\gamma(t))\|_{L^2(Y)}^2 dt.$$

Note that our Corollary 8.18 gives a quantitative version of the quasi-gradient condition $dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}) \geq 0$. A consequence is the following result, which will be useful to us in Section 12 when we establish exponential decay results for trajectories. It says that the energy of an approximate trajectory is commensurable with the drop in F_λ :

Corollary 8.21. *There exists $C_0 > 0$ such that, for any $\lambda \gg 0$ and any closed interval $[t_1, t_2] \subset \mathbb{R}$, the following holds. If $\gamma : [t_1, t_2] \rightarrow W^\lambda \cap B(2R)$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, then*

$$(214) \quad \frac{1}{4C_0} \mathcal{E}(\gamma) \leq F_\lambda(\gamma(t_2)) - F_\lambda(\gamma(t_1)) \leq 4C_0 \mathcal{E}(\gamma).$$

Proof. Note that

$$F_\lambda(\gamma(t_2)) - F_\lambda(\gamma(t_1)) = \int_{t_1}^{t_2} dF_\lambda\left(\frac{d}{dt}\gamma(t)\right)dt = \int_{t_1}^{t_2} dF_\lambda(\mathcal{X}_{q^\lambda}^{\text{gC}}(\gamma(t)))dt.$$

Therefore, the result follows from (213) and Corollary 8.18. \square

¹³Our definition of energy differs from that of [KM07] by a factor of two.

9. GRADINGS

Having established in the previous section that $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient for $\lambda = \lambda_i^\bullet \gg 0$, we are able to define the chain groups of a Morse complex for $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ on $(B(2R) \cap W^\lambda)^\sigma/S^1$ as described in Section 2.6. (After establishing the Morse-Smale condition in Section 10, we will see this indeed gives a complex.) In this section, we relate the gradings of stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ with the gradings of the stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$. An important subtlety here is that stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ can be thought of as living in the infinite-dimensional manifold W^σ/S^1 like those of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ or in the smaller finite-dimensional manifold $(B(2R) \cap W^\lambda)^\sigma/S^1$, and thus there are two ways to define relative gradings. In Section 9.1, we equate these two relative gradings and relate them to the relative gradings for stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$ in an appropriate grading range. In Section 9.2, we define an absolute grading on stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ (a shift of the Morse index) which we relate to the absolute grading $\text{gr}_\mathbb{Q}$ on the stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$.

Thus, the work of Section 7 together with the claimed grading correspondences will establish an identification between the chain groups of the Morse complex for $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ with the monopole Floer complex in an appropriate grading range. This is summarized for the reader's benefit in Section 9.3.

9.1. Relative gradings of stationary points. We would like to define relative gradings on stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ in \mathcal{N}/S^1 using the analogs of the discussions in Section 4.8 and Section 5.11.

First, recall from (167) that we have a section

$$\mathcal{F}_{q^\lambda}^{\text{gC},\tau} : \tilde{\mathcal{C}}^{\text{gC},\tau}(Z) \rightarrow \mathcal{V}^{\text{gC},\tau}(Z).$$

We can take the covariant derivatives of this section, $\mathcal{D}_{\tilde{\gamma}}^{\tilde{g},\tau} \mathcal{F}_{q^\lambda}^{\text{gC},\tau}$ and $\mathcal{D}_{\gamma}^{\tau} \mathcal{F}_{q^\lambda}^{\text{gC},\tau}$, just as in (132) and (135). These agree in the particular case where γ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$. As discussed earlier, it will be easier to work with $\mathcal{D}_{\gamma}^{\tau} \mathcal{F}_{q^\lambda}^{\text{gC},\tau}$, so we focus on this operator.

Lemma 9.1. *For $1 \leq j \leq k$ and $\lambda \gg 0$, the following is true: for any $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$ and each path $[\gamma_\lambda] \in \mathcal{B}_k^{\text{gC},\tau}([x_\lambda], [y_\lambda])$ with representative $\gamma_\lambda \in \mathcal{C}_k^{\text{gC},\tau}(x_\lambda, y_\lambda)$, the operator*

$$(\mathcal{D}_{\gamma}^{\tau} \mathcal{F}_{q^\lambda}^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma_\lambda}^{\text{gC},\tau}} : \mathcal{K}_{j,\gamma_\lambda}^{\text{gC},\tau} \rightarrow \mathcal{V}_{j-1,\gamma_\lambda}^{\text{gC},\tau}$$

is Fredholm, with index equal to $\text{gr}([x_\infty], [y_\infty])$.

Note that here we do not need the assumption that $\lambda = \lambda_i^\bullet$ for the nondegeneracy of $[x_\lambda], [y_\lambda]$, since we have that $[x_\lambda], [y_\lambda]$ are contained in \mathcal{N}/S^1 (cf. Proposition 7.10).

Proof. Recall from Proposition 5.40 that the relative grading between $[x_\infty]$ and $[y_\infty]$ is given by the index of $(\mathcal{D}_{\gamma}^{\tau} \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC},\tau}}$. Recall further from Lemma 5.41 that the index of $(\mathcal{D}_{\gamma}^{\tau} \mathcal{F}_q^{\text{gC},\tau})|_{\mathcal{K}_{j,\gamma}^{\text{gC},\tau}}$ is the same as that of the operator

$$\widehat{Q}_{\gamma}^{\text{gC}} : \mathcal{T}_{j,\gamma}^{\tau} \rightarrow \mathcal{T}_{j-1,\gamma}^{\tau}$$

from (142). We can define a similar operator

$$\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} : \mathcal{T}_{j, \gamma_\lambda}^\tau \rightarrow \mathcal{T}_{j-1, \gamma_\lambda}^\tau,$$

by using the perturbation q^λ instead of q , and the path γ_λ instead of γ . The same arguments as in the proof of Proposition 5.40 show that $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ is Fredholm, so is $(\mathcal{D}_{\gamma_\lambda}^\tau \mathcal{F}_{q^\lambda}^{\text{gC}, \tau})|_{\mathcal{K}_{j, \gamma_\lambda}^{\text{gC}, \tau}}$, and they have the same index. Thus, it remains to show that the operators $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ and $\widehat{Q}_\gamma^{\text{gC}}$ have the same index.

Standard arguments show that the index of the operator $\widehat{Q}_\gamma^{\text{gC}}$ is independent of the choice of γ , for γ a smooth path which is asymptotic to x_∞ and y_∞ . Similarly, it suffices to look at the index of $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ for one choice of γ_λ . We choose γ and γ_λ to be constant outside of the interval $[-T, T]$. As in (145), we write

$$\begin{aligned} \widehat{Q}_\gamma^{\text{gC}} &= \frac{d}{dt} + L_0 + \hat{h}_t^{\text{gC}}, \\ \widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} &= \frac{d}{dt} + L_0 + \hat{h}_t^{\text{gC}, \lambda}. \end{aligned}$$

We will define an intermediate operator $\widehat{Q}^{\text{int}, \lambda}$ interpolating between these two operators. Choose $\beta(t)$ a smooth bump function which is 0 outside of $[-T-2, T+2]$ and is identically 1 in $[-T-1, T+1]$. Define

$$\widehat{Q}^{\text{int}, \lambda} = (1 - \beta(t)) \cdot \widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} + \beta(t) \cdot \widehat{Q}_\gamma^{\text{gC}}.$$

We will compare $\widehat{Q}^{\text{int}, \lambda}$ to $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ and $\widehat{Q}_\gamma^{\text{gC}}$.

We begin by comparing $\widehat{Q}^{\text{int}, \lambda}$ to $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$. We will show that they differ by a compact operator. Notice that

$$\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} - \widehat{Q}^{\text{int}, \lambda} = \beta(t)(\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} - \widehat{Q}_\gamma^{\text{gC}}) = \beta(t)(\hat{h}_t^{\text{gC}, \lambda} - \hat{h}_t^{\text{gC}}).$$

For each t , $\hat{h}_t^{\text{gC}, \lambda} - \hat{h}_t^{\text{gC}}$ is continuous as a self-map from $L_n^2(Y; iT^*Y \oplus \mathbb{S} \oplus i\mathbb{R})$ for each $1 \leq n \leq j$; therefore, this is compact as a map from $L_n^2(Y)$ to $L_{n-1}^2(Y)$. Since $\beta(t)$ is smooth and the terms $\hat{h}_t^{\text{gC}, \lambda}$ and \hat{h}_t^{gC} vary smoothly in t , we have that the four-dimensional map $\beta(t)(\hat{h}_t^{\text{gC}, \lambda} - \hat{h}_t^{\text{gC}})$ from $L_n^2(I \times Y)$ to $L_{n-1}^2(I \times Y)$ is compact whenever I is a compact interval. While the inclusion of $L_n^2(\mathbb{R} \times Y)$ into $L_{n-1}^2(\mathbb{R} \times Y)$ is not compact since $\mathbb{R} \times Y$ is unbounded, we have that $\beta(t)(\hat{h}_t^{\text{gC}, \lambda} - \hat{h}_t^{\text{gC}})$ is compact as an operator from $L_n^2(\mathbb{R} \times Y)$ to $L_{n-1}^2(\mathbb{R} \times Y)$ because β is compactly supported. Therefore, $\widehat{Q}^{\text{int}, \lambda}$ and $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ differ by a compact operator, and hence have the same index.

We now compare $\widehat{Q}^{\text{int}, \lambda}$ to $\widehat{Q}_\gamma^{\text{gC}}$. We write

$$\widehat{Q}^{\text{int}, \lambda} - \widehat{Q}_\gamma^{\text{gC}} = (1 - \beta(t))(\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}} - \widehat{Q}_\gamma^{\text{gC}}).$$

Let us analyze this difference more carefully. When this operator is non-zero (i.e., when $|t| > T+1$), we have that $\widehat{Q}^{\text{int}, \lambda} - \widehat{Q}_\gamma^{\text{gC}}$ is given by $(1 - \beta(t))(\hat{h}_{-\infty}^{\text{gC}, \lambda} - \hat{h}_{-\infty}^{\text{gC}})$ when $t < -T-1$ and $(1 - \beta(t))(\hat{h}_{+\infty}^{\text{gC}, \lambda} - \hat{h}_{+\infty}^{\text{gC}})$ when $t > T+1$. We will argue that $(1 - \beta(t))(\hat{h}_{-\infty}^{\text{gC}, \lambda} - \hat{h}_{-\infty}^{\text{gC}})$, considered as operators from $L_j^2(\mathbb{R} \times Y)$ to $L_j^2(\mathbb{R} \times Y)$, converge to 0 in operator norm as $\lambda \rightarrow \infty$. The same argument will apply for $\hat{h}_{+\infty}^{\text{gC}, \lambda} - \hat{h}_{+\infty}^{\text{gC}}$, and this will imply that $\widehat{Q}^{\text{int}, \lambda} - \widehat{Q}_\gamma^{\text{gC}}$ converges to 0 in operator norm as $\lambda \rightarrow \infty$. In turn, this will imply that for $\lambda \gg 0$, the operators $\widehat{Q}_\gamma^{\text{gC}}$ and $\widehat{Q}^{\text{int}, \lambda}$ must have the same index, completing the proof.

Because γ and γ_λ limit to stationary points of $\mathcal{X}_q^{\text{gC},\sigma}$ and $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ respectively, (146) yields that for $t < -T - 1$,

$$(215) \quad (1 - \beta(t))(\hat{h}_{-\infty}^{\text{gC},\lambda} - \hat{h}_{-\infty}^{\text{gC}}) = (1 - \beta(t))(\widehat{\text{Hess}}_{q^\lambda, x_\lambda}^{\tilde{g}, \sigma} - \widehat{\text{Hess}}_{q, x_\infty}^{\tilde{g}, \sigma})$$

and similarly for $t > T + 1$.

Recall that $\text{Hess}_{q, x_\infty}^\sigma = \mathcal{H}_{x_\infty}^\sigma$ by Lemma 5.26(d). Combining the arguments of Lemma 5.26 with the formula for $\text{Hess}_{q, x_\infty}^\sigma$ in [KM07, p.209], the operator norm (in three-dimensions) of the difference

$$\widehat{\text{Hess}}_{q^\lambda, x_\lambda}^{\tilde{g}, \sigma} - \widehat{\text{Hess}}_{q, x_\infty}^{\tilde{g}, \sigma}$$

from L_n^2 to L_{n-1}^2 can be bounded in terms of the L_n^2 norms of $x_\lambda - x_\infty$, $\mathcal{D}_{\pi(x_\infty)}(\Pi^{\text{elC}} \circ c_q) - \mathcal{D}_{\pi(x_\lambda)}(\Pi^{\text{elC}} \circ p^\lambda c_q)$, and $\mathcal{D}_{x_\infty} \Pi^{\text{elC}} \circ c_q - \mathcal{D}_{x_\lambda} \Pi^{\text{elC}} \circ p^\lambda c_q$, where $\pi(x_\infty), \pi(x_\lambda)$ denote the images in the blow-down. Note that since the two extended Hessians have the same first order term, L_0 , the difference necessarily lands in $L_j^2(Y)$; more generally this difference induces a bounded operator from $L_n^2(Y)$ to $L_n^2(Y)$ for any $1 \leq n \leq j$. For each $1 \leq n \leq j$, the terms mentioned above converge to 0 in $L_n^2(Y)$ as $\lambda \rightarrow \infty$: $x_\lambda - x_\infty$ converges to 0 by the implicit function theorem, while the other two terms converge to zero by combining the arguments of Lemma 6.12 together with the continuity of Π^{elC} and its derivatives (Lemma 3.3). It follows that the norm of

$$A_\lambda := \hat{h}_{-\infty}^{\text{gC},\lambda} - \hat{h}_{-\infty}^{\text{gC}},$$

thought of as an operator from $L_n^2(Y)$ to $L_{n-1}^2(Y)$, converges to zero as $\lambda \rightarrow \infty$. We must extend this to obtain the analogous convergence on $\mathbb{R} \times Y$.

Precisely, we seek to show that the norm of $(1 - \beta(t))A_\lambda$, thought of as an operator from $L_j^2(\mathbb{R} \times Y)$ to $L_{j-1}^2(\mathbb{R} \times Y)$, converges to zero as $\lambda \rightarrow \infty$. Since $\beta(t)$ is smooth and compactly supported, it suffices to consider A_λ itself. Using the analysis on the operator norms from $L_n^2(Y)$ to $L_{n-1}^2(Y)$ above, we have that for $\eta \in L_j^2(\mathbb{R} \times Y)$,

$$\begin{aligned} \|A_\lambda(\eta(t))\|_{L_{j-1}^2(\mathbb{R} \times Y)}^2 &= \sum_{n=0}^{j-1} \int_{\mathbb{R}} \|A_\lambda(\eta^{(n)}(t))\|_{L_{j-n-1}^2(Y)}^2 dt \\ &\leq \sum_{n=0}^j \int_{\mathbb{R}} C_\lambda \|\eta^{(n)}(t)\|_{L_{j-n}^2(Y)}^2 dt \\ &= C_\lambda \|\eta(t)\|_{L_j^2(\mathbb{R} \times Y)}^2, \end{aligned}$$

where $C_\lambda \rightarrow 0$ as $\lambda \rightarrow \infty$. This gives the desired convergence in operator norm and completes the proof. \square

Remark 9.2. We expect that, with more work, one could show that $\widehat{Q}_{\gamma_\lambda, q^\lambda}^{\text{gC}}$ converges to $\widehat{Q}_\gamma^{\text{gC}}$ in operator norm; and hence the two operators have the same index. This would give an alternate proof of the above lemma, without using the intermediate operator $\widehat{Q}^{\text{int}, \lambda}$.

We define a relative grading on the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ in \mathcal{N}/S^1 by

$$(216) \quad \text{gr}([x_\lambda], [y_\lambda]) = \text{ind}(\mathcal{D}_\gamma^\tau \mathcal{F}_{q^\lambda}^{\text{gC}, \tau} |_{\mathcal{K}_{k, \gamma}^{\text{gC}, \tau}}),$$

for any path γ from x_λ to y_λ .

Lemma 9.1 immediately implies the following

Corollary 9.3. *The correspondence $\Xi_\lambda : \mathfrak{C}_\mathcal{N}^\lambda \rightarrow \mathfrak{C}_\mathcal{N}$ from Corollary 7.8 preserves relative gradings.*

Recall from Corollary 7.5 that for $\lambda \gg 0$, the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ contained in \mathcal{N}/S^1 are contained in $(W^\lambda)^\sigma$, so we are able to treat these stationary points as living in finite dimensions. We would like to relate the above grading computations to the Morse indices in finite dimensions. We remind the reader that although $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient, $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ is not a Morse quasi-gradient itself in any obvious way. However, as discussed in Section 2.6, we can still do the analogous constructions in Morse homology.

Proposition 9.4. *Let $[x_\lambda], [y_\lambda] \in \mathcal{N}/S^1$ be stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$. Then, $\text{gr}([x_\lambda], [y_\lambda])$, where this grading is computed in infinite dimensions, is equal to the difference in gradings of $[x_\lambda], [y_\lambda]$ thought of as stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ restricted to $(B(2R) \cap W^\lambda)^\sigma/S^1$.*

Proof. We begin by studying the relevant operators. Let

$$\gamma : \mathbb{R} \rightarrow (B(2R) \cap W^\lambda)^\sigma, \quad \gamma(t) = (a(t), s(t), \phi(t))$$

be a path between stationary points $x_\lambda, y_\lambda \in (B(2R) \cap W^\lambda)^\sigma$. Let also $Z = \mathbb{R} \times Y$.

On the one hand, we have the Fredholm operator

$$(217) \quad Q_\gamma^{\text{gC}} = \mathcal{D}_\gamma^\tau \mathcal{F}_{q^\lambda}^{\text{gC},\tau} \oplus \mathbf{d}_\gamma^{\text{gC},\tau,\ddagger} : \mathcal{T}_{j,\gamma}^{\text{gC},\tau}(x_\lambda, y_\lambda) \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\tau}(Z) \oplus L_{j-1}^2(\mathbb{R}; i\mathbb{R}).$$

This is similar to the operator Q_γ^{gC} from (138), but uses the perturbation q^λ instead of q . (For simplicity, we do not include the perturbation in the notation.)

We now restrict Q_γ^{gC} to

$$(218) \quad \mathcal{T}_{j,\gamma}^{\text{gC},\lambda,\tau} = \{(b, r, \psi) \in \mathcal{T}_{j,\gamma}^{\text{gC},\tau}(x_\lambda, y_\lambda) \mid (b(t), \psi(t)) \in W^\lambda \text{ for all } t\}.$$

After defining $\mathcal{V}_{j-1,\gamma}^{\text{gC},\lambda,\tau}(Z)$ similarly, we get an operator

$$(219) \quad Q_\gamma^{\text{gC},\lambda} = \mathcal{D}_\gamma^\tau \mathcal{F}_{q^\lambda}^{\text{gC},\tau} \oplus \mathbf{d}_\gamma^{\text{gC},\tau,\ddagger} : \mathcal{T}_{j,\gamma}^{\text{gC},\lambda,\tau}(x_\lambda, y_\lambda) \rightarrow \mathcal{V}_{j-1,\gamma}^{\text{gC},\lambda,\tau}(Z) \oplus L_{j-1}^2(\mathbb{R}; i\mathbb{R}).$$

Recall that the index of Q_γ^{gC} computes the relative gradings between stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ in infinite dimensions. We will show below in Lemma 9.5 that Q_γ^{gC} and $Q_\gamma^{\text{gC},\lambda}$ have the same index. Therefore, it remains to show that the index of $Q_\gamma^{\text{gC},\lambda}$ gives the relative gradings in the Morse complex on $(B(2R) \cap W^\lambda)^\sigma/S^1$ in finite dimensions.

Note that the domain of $Q_\gamma^{\text{gC},\lambda}$ consists of paths $(V, \beta) : \mathbb{R} \rightarrow T(W^\lambda)^\sigma \times i\mathbb{R}$, where β comes from the dt component of the connection. Similarly, the codomain also consists of paths (V, β) of this form. We can also view $T_{(a,s,\phi)}(W^\lambda)^\sigma$ as the subspace of $W^\lambda \times \mathbb{R}$ consisting of three-dimensional configurations $((b, \psi), r)$ such that $\text{Re}\langle \phi, \psi \rangle_{L^2(Y)} = 0$.

The first observation is that the L_j^2 norm of V , as a four-dimensional configuration over Z , is equivalent to its L_j^2 norm as a map from \mathbb{R} to the finite-dimensional space $W^\lambda \times \mathbb{R}$. Indeed, in principle the former norm takes into account more derivatives, corresponding to the directions along Y , whereas the latter just uses derivatives in the t direction. However, since we are in W^λ , the j th derivatives of $V(t)$ in the Y directions are bounded (in L^2) by a constant (of the order of λ^j) times the $L^2(Y)$ norm of $V(t)$ itself. This shows that the two Sobolev norms are equivalent. From now on we will use the latter norm for V , i.e. view V as a map from \mathbb{R} to $W^\lambda \times \mathbb{R}$.

By Proposition 7.10, the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ inside \mathcal{N}/S^1 are hyperbolic (where $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ is considered as a vector field in either infinite or finite dimensions). Thus, the relative Morse index between these stationary points is well-defined. We aim to compare $Q_\gamma^{\text{gC},\lambda}$ with the operator

$$(220) \quad \Pi^{\text{agC},\sigma} \circ \left(\frac{D^\sigma}{dt} + \mathcal{D}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC},\sigma} \right) : T_{j,\gamma} \mathcal{P}(x_\lambda, y_\lambda) \rightarrow T_{j-1,\gamma} \mathcal{P}(x_\lambda, y_\lambda).$$

Here, $T_{j,\gamma} \mathcal{P}(x_\lambda, y_\lambda)$ is the subspace of $L_j^2(\mathbb{R}, T(W^\lambda)^\sigma)$ consisting of the paths $V(t) = (b(t), r(t), \psi(t))$ such that $\langle \psi(t), i\phi(t) \rangle_{\tilde{g}} = 0$ for all t , as in (70).

Note that the operator (220) can be used to define the Morse index in $(B(2R) \cap W^\lambda)^\sigma / S^1$. Indeed, it is an operator of the form (2), which uses the connection $\Pi^{\text{agC},\sigma} \circ \mathcal{D}^\sigma$ on the tangent bundle $T((W^\lambda)^\sigma / S^1) \cong \mathcal{K}^{\text{agC},\sigma}$.

Let us write

$$(221) \quad \mathcal{T}_{j,\gamma}^{\text{gC},\lambda,\tau} = T_{j,\gamma} \mathcal{P}(x_\lambda, y_\lambda) \oplus L_j^2(\mathbb{R}; i\mathbb{R}) \oplus L_j^2(\mathbb{R}; i\mathbb{R})$$

where the first $L_j^2(\mathbb{R}; i\mathbb{R})$ summand corresponds to the tangent to the S^1 -orbit (that is, real multiples of $\phi(t)$) and the last summand corresponds to $\beta(t)$. The codomain of $Q_\gamma^{\text{gC},\sigma,\lambda}$ can also be identified with (221), except we use the Sobolev index $j-1$.

With respect to the decomposition $\mathcal{T}_{j,\gamma}^{\text{gC},\lambda,\tau} = \mathcal{V}_{j,\gamma}^{\text{gC},\lambda,\tau} \oplus L_j^2(\mathbb{R}; i\mathbb{R})$, we have

$$(222) \quad Q_\gamma^{\text{gC},\lambda} = \frac{D^\sigma}{dt} + \begin{pmatrix} \mathcal{D}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC},\sigma} & \mathbf{d}_\gamma^{\text{gC},\sigma} \\ \mathbf{d}_\gamma^{\text{gC},\sigma,\tilde{\tau}} & 0 \end{pmatrix},$$

where the operators $\mathbf{d}_\gamma^{\text{gC},\sigma}$ and $\mathbf{d}_\gamma^{\text{gC},\sigma,\tilde{\tau}}$ are the ones defined in (133) and (136). We can further write the entry $\mathcal{D}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ with respect to the decomposition of $T_j(W^\lambda)^\sigma$ into the anticircular Coulomb slice and the tangent to the S^1 -direction, as

$$(223) \quad \begin{pmatrix} \Pi^{\text{agC},\sigma} \circ \mathcal{D}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC},\sigma} & K_2 \\ K_1 & K_3 \end{pmatrix},$$

where the K_i terms are compact and $K_2 = K_3 = 0$ at the endpoints x_λ, y_λ . (Because we use L^2 projections to W^λ in $\mathcal{X}_{q^\lambda}^{\text{gC}}$ instead of \tilde{g} -projections, $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is not a gradient vector field. Therefore, $\mathcal{D} \mathcal{X}_{q^\lambda}^{\text{gC}}$ is not symmetric, and we cannot conclude that the term K_1 vanishes even at stationary points.) However, the tangent to the S^1 -direction depends on time, and hence so does this decomposition of $T_j(W^\lambda)^\sigma$. Therefore, adding $\frac{D^\sigma}{dt}$ to $\mathcal{D}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ does not simply result in adding $\frac{D^\sigma}{dt}$ terms on the diagonal of (223).

Let us identify the tangent to the S^1 -direction with \mathbb{R} , such that the generator $(0, 0, i\phi(t))$ corresponds to 1. We would like to understand the difference

$$(224) \quad \frac{D^\sigma}{dt} - \begin{pmatrix} \Pi^{\text{agC},\sigma} \circ \frac{D^\sigma}{dt} & 0 \\ 0 & \frac{d}{dt} \end{pmatrix}.$$

If we write an element $V \in \mathcal{V}_{j,\gamma}^{\text{gC},\lambda,\tau}$ as a path

$$V(t) = V^{\text{agC}}(t) + f(t) \cdot (0, 0, i\phi(t)),$$

then the expression (224) applied to V becomes

$$(225) \quad \Pi^{\text{gC},\circ} \left(\frac{D^\sigma}{dt} V^{\text{agC}}(t) \right) + f(t) \cdot (0, 0, \Pi_{\phi(t)}^\perp i \frac{d\phi(t)}{dt}),$$

where $\Pi^{\text{gC},\circ}$ is the projection to the span of $(0, 0, i\phi(t))$, with kernel $\mathcal{K}^{\text{agC},\sigma}$. In other words, if $\psi^{\text{agC}}(t)$ is the spinorial component of $V^{\text{agC}}(t)$, then

$$\Pi^{\text{gC},\circ} \left(\frac{D^\sigma}{dt} V^{\text{agC}}(t) \right) = \frac{\left\langle \frac{d}{dt} \psi^{\text{agC}}(t), i\phi(t) \right\rangle_{\tilde{g}}}{\|i\phi(t)\|_{\tilde{g}}^2} \cdot (0, 0, i\phi(t)).$$

By the Leibniz rule, we have

$$\begin{aligned} 0 &= \frac{d}{dt} \langle \psi^{\text{agC}}(t), i\phi(t) \rangle_{\tilde{g}} \\ &= \frac{d}{dt} \text{Re} \langle (0, \psi^{\text{agC}}(t)), \Pi^{\text{elC}}(0, i\phi(t)) \rangle_{L^2(Y)} \\ &= \text{Re} \langle (0, \frac{d}{dt} \psi^{\text{agC}}(t)), \Pi^{\text{elC}}(0, i\phi(t)) \rangle_{L^2(Y)} + \text{Re} \langle (0, \psi^{\text{agC}}(t)), \frac{d}{dt} \Pi^{\text{elC}}(0, i\phi(t)) \rangle_{L^2(Y)} \\ &= \left\langle \frac{d}{dt} \psi^{\text{agC}}(t), i\phi(t) \right\rangle_{\tilde{g}} + \text{Re} \langle (0, \psi^{\text{agC}}(t)), \frac{d}{dt} \Pi^{\text{elC}}(0, i\phi(t)) \rangle_{L^2(Y)}. \end{aligned}$$

Therefore, we can write (225) as

$$- \frac{\text{Re} \langle (0, \psi^{\text{agC}}(t)), \frac{d}{dt} \Pi^{\text{elC}}(0, i\phi(t)) \rangle_{L^2(Y)}}{\|i\phi(t)\|_{\tilde{g}}^2} \cdot (0, 0, i\phi(t)) + f(t) \cdot (0, 0, \Pi_{\phi(t)}^\perp i \frac{d\phi(t)}{dt}).$$

Both of these terms involve the time derivatives of $\phi(t)$, and preserve Sobolev regularity $L_j^2 \rightarrow L_j^2$ as long as $j < k$. We deduce that the difference (224) is compact as a map from L_j^2 to L_{j-1}^2 for $j < k$.

We would like to compare $Q_\gamma^{\text{gC},\lambda}$ to the operator

$$(226) \quad Q_\gamma^{\text{gC},\text{sp},\lambda} = \begin{pmatrix} \Pi^{\text{agC},\sigma} \circ (\frac{D^\sigma}{dt} + \mathcal{D}^\sigma \mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC},\sigma}) & 0 & 0 \\ 0 & \frac{d}{dt} & \mathbf{d}_\gamma^{\text{gC},\sigma} \\ 0 & \mathbf{d}_\gamma^{\text{gC},\sigma,\tilde{\dagger}} & \frac{d}{dt} \end{pmatrix},$$

written with respect to the decomposition (221). We will study the linear interpolations $\alpha Q_\gamma^{\text{gC},\lambda} + (1-\alpha) Q_\gamma^{\text{gC},\text{sp},\lambda}$. With respect to the same decomposition, this linear interpolation can be computed at the endpoints to have the form

$$(227) \quad \begin{pmatrix} \Pi^{\text{agC},\sigma} \circ \mathcal{D}^\sigma \mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC},\sigma} & 0 & 0 \\ \alpha K_1 & 0 & \mathbf{d}_\gamma^{\text{gC},\sigma} \\ 0 & \mathbf{d}_\gamma^{\text{gC},\sigma,\tilde{\dagger}} & 0 \end{pmatrix},$$

where K_1 is as in (223). First, notice that $\Pi^{\text{agC},\sigma} \circ \mathcal{D}^\sigma \mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC},\sigma}$ as a map from $\mathcal{K}_j^{\text{agC},\sigma}$ to $\mathcal{K}_{j-1}^{\text{agC},\sigma}$ is hyperbolic at the endpoints, because the endpoints are hyperbolic stationary points. (It is easy to see that the hyperbolicity for stationary points extends to $j < k$ as well.) We also have that $\begin{pmatrix} 0 & \mathbf{d}_\gamma^{\text{gC},\sigma} \\ \mathbf{d}_\gamma^{\text{gC},\sigma,\tilde{\dagger}} & 0 \end{pmatrix}$ is hyperbolic by the proof of Lemma 5.27. It is straightforward to deduce from here that (227) is hyperbolic at the endpoints, for all α . It follows that

$\alpha Q_\gamma^{\text{gC},\lambda} + (1 - \alpha) Q_\gamma^{\text{gC},\text{sp},\lambda}$ are Fredholm as maps from L_j^2 to L_{j-1}^2 for all α and $j < k$. This implies that they are Fredholm for $j = k$ as well. Indeed, the kernel of the map from L_k^2 to L_{k-1}^2 is necessarily contained in the kernel from L_{k-1}^2 to L_{k-2}^2 and thus is finite-dimensional. A similar argument with the adjoint applies for the cokernel as well. Therefore, $Q_\gamma^{\text{gC},\lambda}$ and $Q_\gamma^{\text{gC},\text{sp},\lambda}$ have the same Fredholm index, being related by a continuous family of Fredholm operators.

For fixed t , both $\mathbf{d}_{\gamma(t)}^{\text{gC},\sigma}$ and $\mathbf{d}_{\gamma(t)}^{\text{gC},\sigma,\tilde{\tau}}$ are just given by multiplication by nonzero constants. Thus, we can write the 2×2 block at the bottom right of (226) as

$$(228) \quad \frac{d}{dt} + A(t),$$

where $A(t)$ is invertible and has real spectrum. Thus, $\{A(t)\}$ has zero spectral flow, and the 2×2 block has index zero. In fact, given the form of $\{A(t)\}$, the operator (228) itself is invertible.

Hence, $Q_\gamma^{\text{gC},\text{sp},\lambda}$ has the same Fredholm index as the top left entry in (226), which is the operator (220). This completes the proof, modulo the claim that Q_γ^{gC} and $Q_\gamma^{\text{gC},\lambda}$ have the same index, which we prove below. \square

Lemma 9.5. *For $\lambda = \lambda_i^\bullet \gg 0$, the index of Q_γ^{gC} is equal to that of $Q_\gamma^{\text{gC},\lambda}$.*

Proof. Since $\lambda = \lambda_i^\bullet$, we have that p^λ is a projection defined with respect to the L^2 (not \tilde{g}) metric.

We begin with some further discussion of the relevant spaces. Define

$$\mathcal{R} = \{(b(t), 0, \psi(t)) \in \mathcal{T}_\gamma^{\text{gC},\tau} \mid (b(t), \psi(t)) \in (W^\lambda)^\perp\}$$

where $(W^\lambda)^\perp$ is the L^2 -orthogonal complement of W^λ in W . In other words, an element of $(W^\lambda)^\perp$ can be decomposed as a (possibly infinite) sum of the eigenvectors of l in W with associated eigenvalue outside of the interval $(-\lambda, \lambda)$. Note that \mathcal{R} does not depend on γ , as $\gamma(t) \in W^\lambda$, so $(b(t), \psi(t)) \in (W^\lambda)^\perp$ automatically implies that $\psi(t)$ is orthogonal to $\gamma(t)$. Because of this, we have a canonical identification of \mathcal{R} with the space of smooth paths in $(W^\lambda)^\perp$, which comes from simply ignoring the middle component. It follows that $\mathcal{R} \oplus \mathcal{T}_\gamma^{\text{gC},\lambda,\tau} = \mathcal{T}_\gamma^{\text{gC},\tau}$. Note that elements of \mathcal{R} have no dt component. We define $\mathcal{R}_j \subset \mathcal{T}_{j,\gamma}^{\text{gC},\tau}$ as the Sobolev completion of \mathcal{R} with respect to the four-dimensional L_j^2 -norm, so that

$$\mathcal{R}_j \oplus \mathcal{T}_{j,\gamma}^{\text{gC},\lambda,\tau} = \mathcal{T}_{j,\gamma}^{\text{gC},\tau}.$$

There is an analogous decomposition

$$\mathcal{V}_{j-1,\gamma}^{\text{gC},\tau} \oplus L_{j-1}^2(\mathbb{R}; i\mathbb{R}) = \mathcal{R}_{j-1} \oplus (\mathcal{V}_{j-1,\gamma}^{\text{gC},\lambda,\tau} \oplus L_{j-1}^2(\mathbb{R}; i\mathbb{R})).$$

With respect to these splittings, we can decompose Q_γ^{gC} as

$$(229) \quad \begin{pmatrix} \Pi_{\mathcal{R}_{j-1}} \circ \mathcal{D}_\gamma^\tau(\mathcal{F}_{q^\lambda}^{\text{gC},\tau})|_{\mathcal{R}_j} & 0 \\ * & Q_\gamma^{\text{gC},\lambda} \end{pmatrix},$$

where the top-right entry of this matrix is zero because $\mathcal{X}_{q^\lambda}^{\text{gC}}$ has image in W^λ , and $\Pi_{\mathcal{R}_{j-1}}$ denotes the L^2 -orthogonal projection to \mathcal{R}_{j-1} . Therefore, to show that $Q_\gamma^{\text{gC},\lambda}$ has the same index as Q_γ^{gC} , it suffices to show that $\Pi_{\mathcal{R}_{j-1}} \circ \mathcal{D}_\gamma^\tau(\mathcal{F}_{q^\lambda}^{\text{gC},\tau})|_{\mathcal{R}_j}$ is invertible. We now compute this operator explicitly.

For $(b(t), 0, \psi(t)) \in \mathcal{R}_j$, it follows from (132) that

$$\begin{aligned} (\mathcal{D}_\gamma^\tau \mathcal{F}_{q^\lambda}^{\text{gC}, \tau})(b(t), 0, \psi(t)) &= \frac{D^\sigma}{dt}(b(t), 0, \psi(t)) + \mathcal{D}_{\gamma(t)}^\sigma \mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}(b(t), 0, \psi(t)) \\ &= \left(\frac{d}{dt}b(t), 0, \Pi_{\phi(t)}^\perp \frac{d}{dt}\psi(t)\right) + \mathcal{D}_{\gamma(t)}^\sigma l^\sigma(b(t), 0, \psi(t)) + \mathcal{D}_{\gamma(t)}^\sigma (p^\lambda c_q)^\sigma(b(t), 0, \psi(t)). \end{aligned}$$

First, note that $\frac{d}{dt}\psi(t)$ and $\phi(t)$ are L^2 -orthogonal. Next, since $p^\lambda c_q$ has image in W^λ , and \mathcal{R}_{j-1} consists of paths of configurations orthogonal to W^λ , it is straightforward to verify that $\Pi_{\mathcal{R}_{j-1}} \circ \mathcal{D}_{\gamma(t)}^\sigma (p^\lambda c_q)^\sigma(b(t), 0, \psi(t)) = 0$. Therefore, we have

$$\Pi_{\mathcal{R}_{j-1}} \circ \mathcal{D}_\gamma^\tau(\mathcal{F}_{q^\lambda}^{\text{gC}, \tau})|_{\mathcal{R}_j} = \frac{d}{dt} + \mathcal{D}_{\gamma(t)}^\sigma l^\sigma.$$

Using (175) and the fact that $\langle \psi(t), D\phi(t) \rangle_{L^2(Y)} = 0$, one can compute directly that

$$(\mathcal{D}_{\gamma(t)}^\sigma l^\sigma)(b(t), 0, \psi(t)) = (*db(t), 0, D\psi(t) - \langle \phi(t), D\phi(t) \rangle_{L^2(Y)} \psi(t)).$$

For notational convenience, let us simply ignore the middle component of $(b(t), 0, \psi(t))$ in \mathcal{R}_j . Define

$$h_t : \mathcal{R}_j \rightarrow \mathcal{R}_j, \quad h_t(b(t), \psi(t)) := (0, -\langle \phi(t), D\phi(t) \rangle_{L^2(Y)} \psi(t)).$$

Showing that $\Pi_{\mathcal{R}_{j-1}} \circ \mathcal{D}_\gamma^\tau(\mathcal{F}_{q^\lambda}^{\text{gC}, \tau})|_{\mathcal{R}_j}$ is invertible is equivalent to showing the invertibility of

$$\frac{d}{dt} + l + h_t : \mathcal{R}_j \rightarrow \mathcal{R}_{j-1}.$$

First, we prove that $\frac{d}{dt} + l + h_t$ is injective. Since $\gamma(t) \in (W^\lambda)^\sigma$ for all t and $\|\phi(t)\|_{L^2(Y)} = 1$ for all t , it follows that there exists $\epsilon > 0$ such that $|\langle \phi(t), D\phi(t) \rangle_{L^2(Y)}| \leq \lambda - \epsilon$, independent of t . Suppose that $(b(t), \psi(t))$ is in the kernel of $\frac{d}{dt} + l + h_t$ and write $(b(t), \psi(t)) = \sum_{\kappa \geq \lambda} (b_\kappa(t), \psi_\kappa(t))$, where we are summing according to the eigenspace decomposition of l . Note that $(b_\kappa(t), \psi_\kappa(t))$ is in the kernel of $\frac{d}{dt} + l + h_t$ for each κ . However, it is straightforward to verify as in the proofs of Lemmas 6.14 and 6.15, that since $\kappa \geq \lambda$ we must have that $(b_\kappa(t), \psi_\kappa(t))$ must be unbounded either as $t \rightarrow \infty$ or $t \rightarrow -\infty$. This contradicts $(b_\kappa(t), \psi_\kappa(t))$ being an L_j^2 -path.

It remains to see that $\frac{d}{dt} + l + h_t$ is surjective. Note that $\frac{d}{dt} + l + h_t$ naturally extends to an operator on sections from $\mathbb{R} \times Y$ to $p^*(iT^*Y \oplus \mathbb{S})$. The formal adjoint of this operator is $-\frac{d}{dt} + l + h_t$, which is injective by the same argument as above. Therefore, the extension of $\frac{d}{dt} + l + h_t$ is surjective. Since the formal adjoint preserves the condition of paths being in $(W^\lambda)^\perp$, we see that $\frac{d}{dt} + l + h_t$, as defined on \mathcal{R}_j , must be surjective. \square

So far we have only discussed the relative grading between stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ that live in \mathcal{N}/S^1 . Let us end with a discussion about the reducible stationary points that are in $(B(2R) \cap W^\lambda)^\sigma/S^1$, but not necessarily in \mathcal{N}/S^1 .

For the rest of the subsection we fix $\lambda = \lambda_i^\bullet$ sufficiently large, and a reducible stationary point $(a, 0)$ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ in $B(2R)$. Consider the reducible stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ inside $(B(2R) \cap W^\lambda)^\sigma/S^1$ that are of the form $[(a, 0, \phi)]$. We write $\kappa(\phi)$ for the associated eigenvalues. By Proposition 7.11, any such $[(a, 0, \phi)]$ is hyperbolic, when thought of as a stationary

point on the finite-dimensional manifold $(B(2R) \cap W^\lambda)^\sigma / S^1$. Since $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient, we can compute the relative gradings (in finite dimensions) between these points.

Lemma 9.6. *Let $[(a, 0, \phi)]$ and $[(a, 0, \phi')]$ be stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ as above. Assume that $\kappa(\phi) > \kappa(\phi')$. Then, the relative grading between these points, as computed from $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ restricted to the finite-dimensional manifold $(B(2R) \cap W^\lambda)^\sigma / S^1$, is given by*

(230)

$$\text{gr}([(a, 0, \phi)], [(a, 0, \phi')]) = \begin{cases} 2i(\kappa(\phi), \kappa(\phi')) & \text{if } \kappa(\phi) \text{ and } \kappa(\phi') \text{ have the same sign,} \\ 2i(\kappa(\phi), \kappa(\phi')) - 1 & \text{otherwise.} \end{cases}$$

Proof. This follows from Lemma 2.19. \square

Lemma 9.7. *Suppose $[(a, 0, \phi)]$ is a stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ that is contained in $(B(2R) \cap W^\lambda)^\sigma / S^1$ but not in \mathcal{N} / S^1 .*

(a) *If $\kappa(\phi) > 0$, then for all stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ of the form $[(a, 0, \phi')]$ that are contained in \mathcal{N} / S^1 , we have $\text{gr}([(a, 0, \phi)], [(a, 0, \phi')]) \geq 2$.*

(b) *If $\kappa(\phi) < 0$, then for all stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ of the form $[(a, 0, \phi')]$ that are contained in \mathcal{N} / S^1 , we have $\text{gr}([(a, 0, \phi)], [(a, 0, \phi')]) \leq -2$.*

Proof. Consider a pair of reducible stationary points $[x]$ and $[y]$ of $\mathcal{X}_q^{\text{agC}, \sigma}$ in \mathcal{N} / S^1 with the same connection component. It follows from Corollary 9.3, Proposition 9.4, and Lemma 9.6, that the relative grading between $[x]$ and $[y]$ is the same as the relative grading between $[x_\lambda]$ and $[y_\lambda]$, considered as stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ on $(B(2R) \cap W^\lambda)^\sigma / S^1$. Further, the spinorial energies of $[x]$ and $[y]$ are necessarily close to the λ -spinorial energies of $[x_\lambda]$ and $[y_\lambda]$ respectively. Recall that $[x_\lambda]$ and $[y_\lambda]$ are necessarily contained in \mathcal{N} / S^1 . Equation (58) and Lemma 9.6 give that in each case, the relative gradings are computed in terms of the orderings by eigenvalues, which correspond to (λ) -spinorial energy. In particular, this implies that if $[(a, 0, \phi)]$ is a reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ not in \mathcal{N} / S^1 , then its associated eigenvalue cannot sit between those of $[x_\lambda]$ and $[y_\lambda]$ for any pair $[x_\lambda]$ and $[y_\lambda]$. The result now follows. \square

The following is an immediate consequence of the proof of Lemma 9.7.

Corollary 9.8. *Let $[x]$ denote the reducible stationary point of $\mathcal{X}_q^{\text{agC}, \sigma}$ which has lowest eigenvalue among all reducible stationary points with the same connection component. Then, $[x_\lambda]$ is the reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ with the lowest positive eigenvalue among those in $(B(2R) \cap W^\lambda)^\sigma / S^1$ with the same connection component.*

9.2. Absolute gradings. Recall that Theorem 1.1 asserts an isomorphism of $\widetilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s}))$ with $\widetilde{HM}(Y, \mathfrak{s})$ which respects the absolute gradings. As our current strategy for the proof of this isomorphism is to identify each of these modules in a certain grading range with the Morse homology of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ on $(B(2R) \cap W^\lambda)^\sigma / S^1$, we need to define an absolute grading on the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ which lines up with the gradings coming from SWF and from \widetilde{HM} .

In Section 8, we showed that $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient. From this, (22) implies that the Morse complex for $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ computes the reduced S^1 -equivariant homology of the Conley index I^λ from Section 3 in a certain grading range. Since $\text{SWF}(Y, \mathfrak{s}) = \Sigma^{-n(Y, \mathfrak{s}, g)} \mathbb{C} \Sigma^{-W^{(-\lambda, 0)}} I^\lambda$, to have a complex whose homology agrees with that of the reduced S^1 -equivariant homology of the spectrum, we must shift the gradings accordingly. Therefore, for $[x_\lambda]$ a stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, define

$$(231) \quad \text{gr}_\lambda^{\text{SWF}}([x_\lambda]) := \text{ind}([x_\lambda] \text{ in } (W^\lambda)^\sigma / S^1) - \dim W^{(-\lambda, 0)} - 2n(Y, \mathfrak{s}, g),$$

where $n(Y, \mathfrak{s}, g)$ is the quantity mentioned at the end of Section 3. Therefore, the Morse complex for $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ with absolute grading given instead by $\text{gr}_\lambda^{\text{SWF}}$ computes $\tilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s}))$ in the appropriate grading range by the discussion in Section 2.8.

Thus, to connect the gradings on the Floer spectrum with monopole Floer homology, we will need to relate the absolute grading $\text{gr}_\lambda^{\text{SWF}}$ on \mathfrak{C}_λ with the absolute grading $\text{gr}^\mathbb{Q}$ from (59) defined on \mathfrak{C}_λ . This is the subject of the following proposition.

Proposition 9.9. *For any $\lambda = \lambda_i^\bullet \gg 0$ and $[x] \in \mathfrak{C}$, we have*

$$(232) \quad \text{gr}_\lambda^{\text{SWF}}([x_\lambda]) = \text{gr}^\mathbb{Q}([x]).$$

In order to prove (232), let us now recall the precise definition of $n(Y, \mathfrak{s}, g)$ from [Man03, Equation (6)]:

$$(233) \quad n(Y, \mathfrak{s}, g) = \text{ind}_\mathbb{C}(D^+) - \frac{c_1(\mathfrak{t})^2 - \sigma(X)}{8}.$$

Here: X is a simply connected, oriented, compact Riemannian four-manifold with boundary Y (such that the metric is a product near the boundary); \mathfrak{t} is a Spin^c structure on X such that $\mathfrak{t}|_Y = \mathfrak{s}$; D^+ is the Dirac operator on (X, \mathfrak{t}) with Atiyah-Patodi-Singer boundary conditions, and associated to a connection extending A_0 on Y ; and $\text{ind}_\mathbb{C}$ denotes the index of a complex operator, which is twice the real index $\text{ind}_\mathbb{R}$. The Atiyah-Patodi-Singer boundary conditions mean that the domain of D^+ consists of spinors whose restrictions to Y project trivially to the negative eigenspaces of the three-dimensional Dirac operator D .

Proof of Proposition 9.9. Let $[x] \in \mathfrak{C}^s$ be a reducible generator corresponding to the lowest positive eigenvalue of the operator $D_{q, a}$, where $[x] = [(a, 0, \phi)]$. Corollary 9.3 and Proposition 9.4 imply that

$$(234) \quad \text{gr}_\lambda^{\text{SWF}}([x_\lambda]) - \text{gr}_\lambda^{\text{SWF}}([y_\lambda]) = \text{gr}([x_\infty], [y_\infty]),$$

so it suffices to prove the relation for $[x]$.

Pick (X, \mathfrak{t}) as in the definition of $n(Y, \mathfrak{s}, g)$ above. Then, recall from Section 4.8 that

$$\text{gr}^\mathbb{Q}([x]) = -\text{gr}_z(X, [x]) + \frac{c_1(\mathfrak{t})^2 - \sigma(X)}{4} - b^+(X) - 1.$$

We seek to show that

$$(235) \quad \text{gr}_z(X, [x]) = \text{ind}_\mathbb{R}(D^+) - b^+(X) - 1 - \text{ind}([x_\lambda] \text{ in } (W^\lambda)^\sigma / S^1) + \dim W^{(-\lambda, 0)}.$$

The quantity $\text{gr}_z(X, [x])$ is the virtual dimension of the Seiberg-Witten moduli space on X (with an added cylindrical end) with asymptotics given by x . Following the proof of Lemma 28.3.2 in [KM07], we can compute $\text{gr}_z(X, [x])$ by using a reducible configuration on X . It then becomes the index of an operator with two parts: one is a perturbed signature operator, and the other is a perturbed Dirac operator. The former would have index

$-b^+(X) - 1$ if the perturbation \mathfrak{q} were zero, but in general it differs from this by the index of a signature operator on the cylinder $[0, 1] \times Y$, with boundary data $(0, 0)$ and (\mathfrak{q}, a) . This index can be computed as the spectral flow of the family

$$\begin{pmatrix} 0 & -d^* \\ -d & *d + 2t\mathcal{D}_{(ta,0)}\mathfrak{q}^0 \end{pmatrix} : \Omega^0(Y; i\mathbb{R}) \oplus \Omega^1(Y; i\mathbb{R}) \rightarrow \Omega^0(Y; i\mathbb{R}) \oplus \Omega^1(Y; i\mathbb{R}), \quad t \in [0, 1].$$

By a compact perturbation that keeps the endpoints fixed, we can change this family of operators into

$$(236) \quad \begin{pmatrix} 0 & -d^* \\ -d & *d + 2t\mathcal{D}_{(a,0)}\mathfrak{q}^0 \end{pmatrix}, \quad t \in [0, 1].$$

Since $(a, 0)$ is a stationary point, we have that

$$\mathcal{D}_{(a,0)}\mathfrak{q}^0 = \begin{pmatrix} 0 & 0 \\ 0 & \mathcal{D}_{(a,0)}\eta_{\mathfrak{q}}^0 \end{pmatrix}$$

with respect to the decomposition of imaginary one-forms into $\ker d \oplus \ker d^*$. Hence, (236) decomposes into a 3×3 block form, where one block is constantly $\begin{pmatrix} 0 & -d^* \\ -d & 0 \end{pmatrix}$ and the other is

$$*d + 2t\mathcal{D}_{(a,0)}\eta_{\mathfrak{q}}^0 : \ker d^* \rightarrow \ker d^*, \quad t \in [0, 1].$$

Since $\begin{pmatrix} 0 & -d^* \\ -d & 0 \end{pmatrix}$ has no spectral flow, we have reduced the computation to the spectral flow of this last family. Let us denote the spectral flow by $\text{SF}(\mathfrak{q})^0$.

The second contribution to $\text{gr}_z(X, [x])$ comes from the index of a perturbed Dirac operator $D_{\mathfrak{q},a}^+ - \lambda_0$, where $D_{\mathfrak{q},a}^+$ is an APS operator but with $D_{\mathfrak{q},a}$ on the boundary, unlike $D^+ = D_{0,0}^+$. Here λ_0 is the eigenvalue corresponding to x , and the domain of $D_{\mathfrak{q},a}^+ - \lambda_0$ consists of spinors whose restrictions to Y project trivially to the eigenspaces of $D_{\mathfrak{q},a}$ for eigenvalues $< \lambda_0$. Since λ_0 is the lowest positive eigenvalue, the domain is the same as the one we considered for D^+ in (233). The two operators $D_{\mathfrak{q},a}^+$ and $D_{\mathfrak{q},a}^+ - \lambda_0$ differ by a constant (hence compact) term, and hence have the same index. Note that

$$\text{ind}_{\mathbb{R}}(D_{\mathfrak{q},a}^+) - \text{ind}_{\mathbb{R}}(D^+) = \text{SF}(\mathfrak{q})^1,$$

where $\text{SF}(\mathfrak{q})^1$ is the spectral flow of the perturbed Dirac operators on Y as we move from $(0, 0)$ to (\mathfrak{q}, a) . Note that at a reducible stationary point $(a, 0)$, we have that $D_a\psi + \mathcal{D}_{(a,0)}\mathfrak{q}^1(0, \psi) = D\psi + \mathcal{D}_{(a,0)}(c_{\mathfrak{q}})^1(0, \psi)$, or in short, $D_{\mathfrak{q},a}^{\text{gC}} = D_{\mathfrak{q},a}$. In particular, we can compute this spectral flow in Coulomb gauge.

Therefore, we have

$$\text{gr}_z(X, [x]) = \text{ind}_{\mathbb{R}}(D^+) - b^+(X) - 1 + \text{SF}(\mathfrak{q})^0 + \text{SF}(\mathfrak{q})^1.$$

To obtain (235), it remains to show that

$$(237) \quad \dim W^{(-\lambda,0)} - \text{ind}([x_{\lambda}] \text{ in } (W^{\lambda})^{\sigma}/S^1) = \text{SF}(\mathfrak{q})^0 + \text{SF}(\mathfrak{q})^1.$$

We now analyze the terms on the left-hand side of (237). The first term, $\dim W^{(-\lambda,0)}$, is the number of negative eigenvalues of l (counted with multiplicity) between $-\lambda$ and 0. The second term requires a more careful analysis. Write $[x_{\lambda}] = (a_{\lambda}, 0, \phi_{\lambda})$. By Corollary 9.8, $[x_{\lambda}]$ has lowest positive eigenvalue among stationary points in $(B(2R) \cap W^{\lambda})^{\sigma}/S^1$ of $\mathcal{X}_{\mathfrak{q}_{\lambda}}^{\text{agC},\sigma}$ with connection component a_{λ} (and not just among those in \mathcal{N}/S^1). By Lemma 2.19, $\text{ind}([x_{\lambda}] \text{ in } (W^{\lambda})^{\sigma}/S^1)$ is the sum of two parts. The first is the number of negative eigenvalues of the linearization of $l + p^{\lambda}c_{\mathfrak{q}}$ restricted to the connection summand of W^{λ} , i.e.

$*d + \mathcal{D}_{(a_\lambda, 0)}(p^\lambda c_q)^1(\cdot, 0)$. The second part is the number of negative eigenvalues of the linearization of $l + p^\lambda c_q$, restricted to the spinorial summand of W^λ , at $(a_\lambda, 0)$, i.e. $D + \mathcal{D}_{(a_\lambda, 0)}(p^\lambda c_q)^1(0, \cdot)$.

Putting it all together, we find that the left-hand side of (237) is the spectral flow from l to $l + p^\lambda A_\lambda$, where

$$A_\lambda(b, \psi) = (\mathcal{D}_{(a_\lambda, 0)} c_q^0(b, 0), \mathcal{D}_{(a_\lambda, 0)} c_q^1(0, \psi)),$$

when considered as operators on W^λ . However, since λ is of the form λ_i^\bullet , this is the same as the spectral flow from l to $l + p^\lambda A_\lambda p^\lambda$, considered as operators from W_k to W_{k-1} . Thus, to establish (237), it remains to show that there is no spectral flow from $l + p^\lambda A_\lambda p^\lambda$ to $l + A_\infty$, as operators from W_k to W_{k-1} . For the rest of the discussion, we will only consider operators from W_k to W_{k-1} . Further, all these operators will have index zero (being compact perturbations of l), and hence for them injectivity or surjectivity is equivalent to invertibility. Due to the block form of these operators, this fact remains true if we restrict the operators to either their connection or spinorial components.

Since $[x_\infty]$ is a non-degenerate stationary point of $\mathcal{X}_q^{\text{agC}, \sigma}$, we have that $l + A_\infty$ is injective. (This follows from Proposition 4.13 and that this operator is index 0.) By Proposition 7.11, we have that $[x_\lambda]$ is a non-degenerate reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, and from this it follows that $l + p^\lambda A_\lambda p^\lambda$ is injective. Since A_∞ and $p^\lambda A_\lambda p^\lambda$ are L^2 self-adjoint, it suffices to show that for $\lambda \gg 0$ and $t \in [0, 1]$, $l + h_t^\lambda$ is injective, where h_t^λ is the compact operator

$$h_t^\lambda = t p^\lambda A_\lambda p^\lambda + (1 - t) A_\infty.$$

Note that for any sequence v_n which converges to v weakly in W_k , any sequence $t_n \in [0, 1]$, and $\lambda_n \rightarrow \infty$, we have

$$h_{t_n}^{\lambda_n}(v_n) \rightarrow A_\infty(v) \text{ in } W_{k-1}.$$

Now suppose that $l + h_{t_n}^{\lambda_n}$ is not injective for some sequences $t_n \in \mathbb{R}$ and $\lambda_n \rightarrow \infty$. Let $v_n \in W_k$ with $\|v_n\|_{L_k^2} = 1$ be such that $l + h_{t_n}^{\lambda_n}(v_n) = 0$. The v_n converge weakly in W_k to some v and as discussed, $h_{t_n}^{\lambda_n}(v_n)$ converges in W_{k-1} to $A_\infty(v)$. Thus, we see that $l(v_n)$ converges in W_{k-1} to $l(v)$. In particular, v_n converges in W_k to v , which consequently has $\|v\|_{L_k^2} = 1$, and $(l + A_\infty)(v) = 0$. This contradicts the injectivity of $l + A_\infty$. The relation (237) follows. \square

Having established Proposition 9.9, we can rephrase Lemma 9.7 as the following.

Proposition 9.10. *Any reducible stationary point of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ in $(B(2R) \cap W^\lambda)^\sigma / S^1$ which is in the grading range $[-N, N]$ is contained in \mathcal{N} / S^1 , assuming $\lambda = \lambda_i^\bullet$ for some $i \gg 0$.*

9.3. Conclusions. Recall, from the discussion at the beginning of Section 7.4, that approximate irreducible stationary points are necessarily in \mathcal{N} / S^1 . Therefore, using the fact that Ξ_λ is grading preserving and the fact that irreducible stationary points have vanishing (λ) -spinorial energy, it follows that there exists $N > 0$ such that all irreducible stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ have grading in $[-N, N]$ for all $\lambda = \lambda_i^\bullet \gg 0$. Here, recall that we grade the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ using $\text{gr}_\lambda^{\text{SWF}}$, defined in (231). We can now summarize the results of this section in the following.

Proposition 9.11. *Let \mathfrak{q} be a very tame, admissible perturbation, and fix $N > 0$ such that all irreducible stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$ have grading in $[-N, N]$ for all $\lambda = \lambda_i^\bullet \gg 0$. For all $\lambda = \lambda_i^\bullet \gg 0$, there is a one-to-one correspondence Ξ_λ between:*

- *the stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$ in $(B(2R) \cap W^\lambda)^\sigma/S^1$ with grading in $[-N, N]$, including all irreducibles, and*
- *the stationary points of $\mathcal{X}_{\mathfrak{q}}^{\text{agC},\sigma}$ with grading in $[-N, N]$, including all irreducibles.*

This correspondence preserves the grading, as well as the type of stationary point (irreducible, stable, unstable). Furthermore, all the stationary points of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{agC},\sigma}$ in $(B(2R) \cap W^\lambda)^\sigma/S^1$ with grading in $[-N, N]$ are hyperbolic.

Proof. The conclusion follows from Corollary 7.8, Proposition 7.10, Corollary 9.3 and Proposition 9.10. \square

10. THE MORSE-SMALE CONDITION FOR THE APPROXIMATE FLOW

Recall that in Section 8, we established that $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}}$ is a Morse equivariant quasi-gradient vector field on $W^\lambda \cap B(2R)$. In this section, we show that it is also Morse-Smale, in the sense of Definition 2.24. Recall, from Lemma 2.2 and the discussion at the end of Section 2.8 that the Morse-Smale condition can be rephrased in terms of the surjectivity of a linear operator. In our setting, let

$$\gamma : \mathbb{R} \rightarrow (W^\lambda \cap B(2R))^\sigma \subset (W^\lambda)^\sigma,$$

be a trajectory of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$, going between two stationary points x and y . Regularity of the moduli space at $[\gamma]$ is equivalent to the surjectivity of the operator

$$(238) \quad \Pi^{\text{agC},\sigma} \circ \left(\frac{D^\sigma}{dt} + \mathcal{D}^\sigma \mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma} \right) : T_{j,\gamma} \mathcal{P}(x, y) \rightarrow T_{j-1,\gamma} \mathcal{P}(x, y)$$

which has already made an appearance in (220). Again, while it may seem like a more natural choice to work with an operator where the derivatives are given by a connection coming from the \tilde{g} -metric, the choice of connection does not matter at a trajectory.

Alternatively, we can view γ as a path in the infinite dimensional space W^σ . The corresponding linearized operator is

$$(239) \quad \mathcal{D}_\gamma^\tau \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau} : \mathcal{T}_{k,\gamma}^{\text{gC},\tau}(x, y) \rightarrow \mathcal{V}_{k-1,\gamma}^{\text{gC},\tau}(Z).$$

Lemma 10.1. *The operator (238) is surjective if and only if the operator (239) is surjective.*

Proof. We use the notation and the results from Lemma 5.41 and the proof of Proposition 9.4. By the analogue of Lemma 5.41 with \mathfrak{q}^λ instead of \mathfrak{q} , surjectivity of (239) is equivalent to that of the operator Q_γ^{gC} from (217). In the proof of Lemma 9.5 we gave a block form (229) for Q_γ^{gC} , and we also showed that the top left block is invertible. This implies that Q_γ^{gC} is surjective if and only if the operator $Q_\gamma^{\text{gC},\lambda}$ from (219) is surjective. Finally, $Q_\gamma^{\text{gC},\lambda}$ can be related to (238) through (226), using the invertibility of the operator (228). \square

Thus, we can work with the operators $\mathcal{D}_\gamma^\tau \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau}$. Note that we should only ask for their surjectivity when γ is boundary-unobstructed. When γ is boundary-obstructed, we require surjectivity of its summand $(\mathcal{D}_\gamma^\tau \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau})^\partial$, which acts on the spaces of paths in the boundary of $(W^\lambda)^\sigma$. Similar arguments as above apply for the boundary-obstructed case as well.

Proposition 10.2. *We can choose the admissible perturbation \mathfrak{q} such that for any $\lambda \in \{\lambda_1^\bullet, \lambda_2^\bullet, \dots\}$ sufficiently large, the following holds. Given any flow trajectory γ for the restriction of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$ to $(B(2R) \cap W^\lambda)^\sigma$, we have that:*

- (i) *If γ is boundary-unobstructed, then $\mathcal{D}_\gamma^\tau \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau}$ is surjective;*
- (ii) *If γ is boundary-obstructed, then $(\mathcal{D}_\gamma^\tau \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau})^\partial$ is surjective.*

Proof. This is similar to the proof of transversality for moduli spaces of trajectories in [KM07, Section 15]. So far, we have chosen an admissible perturbation \mathfrak{q}_0 so that the stationary points of both $\mathcal{X}_{\mathfrak{q}_0}^{\text{gC},\sigma}$ and $\mathcal{X}_{\mathfrak{q}_0^\lambda}^{\text{gC},\sigma}$ inside $B(2R)^\sigma$ are non-degenerate; cf. Proposition 7.11. Consider the blow-down projections of the stationary points of $\mathcal{X}_{\mathfrak{q}_0}^{\text{gC},\sigma}$; these come in a finite number of S^1 -orbits. Pick disjoint open neighborhoods of those orbits, and

let U be the union of these neighborhoods. By Lemma 7.9, for λ large, the blow-down projections of the stationary points of $\mathcal{X}_{\mathbf{q}_0}^{\text{gC},\sigma}$ from $B(2R)^\sigma$ land inside U . Consider the set of perturbations

$$\mathcal{P}_U = \{\mathbf{q} \in \mathcal{P} \mid \mathbf{q}|_U = \mathbf{q}_0|_U\}.$$

We can find an open neighborhood $\nu(\mathbf{q}_0)$ of \mathbf{q}_0 in \mathcal{P}_U such that for all \mathbf{q} in this neighborhood, we have the same set of stationary points for $\mathcal{X}_{\mathbf{q}}^{\text{gC},\sigma}$ and $\mathcal{X}_{\mathbf{q}_0}^{\text{gC},\sigma}$ as for \mathbf{q}_0 , and therefore we still have nondegeneracy for them.

We now claim that for a residual set of perturbations \mathbf{q} in $\nu(\mathbf{q}_0) \subset \mathcal{P}_U$, the desired surjectivity conditions hold. Since we work with a countable set of λ , it suffices to prove this for some fixed λ , sufficiently large. We define a parametrized map

$$\mathfrak{M} : \tilde{\mathcal{C}}_k^{\text{gC},\tau}(x, y) \times \nu(\mathbf{q}_0) \rightarrow \mathcal{V}_{k-1}^{\text{gC},\tau}(Z), \quad \mathfrak{M}(\gamma, \mathbf{q}) = \mathcal{F}_{\mathbf{q}_\lambda}^{\text{gC},\tau}(\gamma).$$

When γ is reducible (i.e., contained in the reducible locus), there is a similar map $(\mathfrak{M})^\partial$ acting on the space of paths in the boundary.

To prove our claim, it is enough to check that the derivative of \mathfrak{M} is surjective at all points (γ, \mathbf{q}) in $\mathfrak{M}^{-1}(0)$ when γ is irreducible; and that the derivative of $(\mathfrak{M})^\partial$ is surjective at reducibles. The proof of these facts is entirely similar to the proof of Proposition 15.1.3 in [KM07]. \square

Let us now put the results of Sections 7, 8, 9, and the current one in context. Recall that it is our goal to establish an isomorphism between $\widetilde{HM}(Y, \mathfrak{s})$ and the (singular) equivariant homology $\tilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s}))$. To do this, for each $N \gg 0$, we will establish an isomorphism between the truncations $\widetilde{HM}_{\leq N}(Y, \mathfrak{s})$ and $\tilde{H}_{\leq N}^{S^1}(\text{SWF}(Y, \mathfrak{s}))$ by passing through an intermediate group, the S^1 -equivariant Morse homology of $\mathcal{X}_{\mathbf{q}_\lambda}^{\text{agC},\sigma}$ on $(B(2R) \cap W^\lambda)^\sigma / S^1$ as defined in Section 2. This group can be defined by Propositions 8.1 and 10.2 for $\lambda = \lambda_i^\bullet$ with $i \gg 0$. We have established in Section 2 that the homology of this Morse complex in an appropriate grading range will be isomorphic to the (singular) Borel homology $\tilde{H}_{\leq N}^{S^1}(\text{SWF}(Y, \mathfrak{s}))$. Thus, it remains to identify the S^1 -equivariant Morse homology of $\mathcal{X}_{\mathbf{q}_\lambda}^{\text{agC},\sigma}$ with monopole Floer homology in the corresponding grading range.

At this point, for $\lambda = \lambda_i^\bullet$ with $i \gg 0$, we have established in Proposition 9.11 a correspondence between the stationary points of $\mathcal{X}_{\mathbf{q}}^{\text{agC},\sigma}$ and $\mathcal{X}_{\mathbf{q}_\lambda}^{\text{agC},\sigma}$ with grading in the interval $[-N, N]$. Thus, using the results of Section 5, we have an isomorphism on the level of graded chain groups (but not yet on homology) between $\widetilde{CM}_{\leq N}(Y, \mathfrak{s}, \mathbf{q})$ and the S^1 -equivariant Morse complex for $\mathcal{X}_{\mathbf{q}_\lambda}^{\text{agC},\sigma}$.

This leaves us with one major step, which is to construct a chain complex isomorphism between the S^1 -equivariant Morse complex and $\widetilde{CM}_{\leq N}(Y, \mathfrak{s}, \mathbf{q})$ by relating the trajectories of $\mathcal{X}_{\mathbf{q}}^{\text{agC},\sigma}$ to those in finite dimensions; this will be analogous to the correspondence on the level of stationary points we have established in this section. This is the focus of Section 12 and Section 13. Before doing so, in the next section we do some technical work which will allow us to relate paths between stationary points of $\mathcal{X}_{\mathbf{q}_\lambda}^{\text{gC},\sigma}$ to paths between stationary points of $\mathcal{X}_{\mathbf{q}}^{\text{gC},\sigma}$.

11. SELF-DIFFEOMORPHISMS OF CONFIGURATION SPACES

In Section 7.2, we established a correspondence $\Xi_\lambda : \mathfrak{C}_\mathcal{N}^\lambda \rightarrow \mathfrak{C}_\mathcal{N}$ between stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ and $\mathcal{X}_q^{\text{agC},\sigma}$. Our goal is to be able to do this for trajectories as well. In this case, the trajectories live in different path spaces: trajectories of $\mathcal{X}_q^{\text{agC},\sigma}$ live in $\mathcal{B}_k^{\text{gC},\tau}([x_\infty], [y_\infty])$ while we will see in Section 12 that trajectories of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ are in $\mathcal{B}_k^{\text{gC},\tau}([x_\lambda], [y_\lambda])$. Therefore, we need a way to relate these different spaces. In this section, we will extend the correspondence from Corollary 7.8 first to a family of S^1 -equivariant self-diffeomorphisms of W^σ , and then to self-diffeomorphisms of $W^\tau(I \times Y)$ for $I \subset \mathbb{R}$ and to other path spaces. This will be needed in Proposition 12.22 and Proposition 13.2. The construction of these maps will use a setup similar to that of the function T_λ defined in Section 8.2.

Before stating the first result, we need some preliminaries. In this section, in order to make statements about the smoothness of functions with respect to λ as a parameter, we do *not* restrict to the case that $\lambda = \lambda_i^\bullet$. We will not use the results of Section 8 or Section 9, so this will not be a problem. For each stationary point x_∞ of $\mathcal{X}_q^{\text{gC},\sigma}$ in \mathcal{N} let x_λ denote the stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ which is L^2 closest to x_∞ . (For an explanation of the well-definedness of x_λ see Section 8.2.) In what follows, we will abuse notation and say that a function G_λ , which depends on λ , is *smooth in λ at and near infinity* if there exists $\epsilon > 0$ such that $G_{f^{-1}(r)}$ depends smoothly on $r \in [0, \epsilon)$, where $f : (0, \infty] \rightarrow [0, 1)$ is the homeomorphism from Section 7.2.

Lemma 11.1. *For $\lambda \gg 0$, there exists an S^1 -equivariant diffeomorphism $\Xi_\lambda : W_0^\sigma \rightarrow W_0^\sigma$ satisfying:*

- (i) Ξ_λ sends x_λ to x_∞ for each stationary point $x_\infty \in \mathcal{N}$,
- (ii) Ξ_λ restricts to a self-diffeomorphism of W_j^σ for any $1 \leq j \leq k$,
- (iii) Let Ξ_∞ be the identity. Then, for $0 \leq j \leq k$, $\Xi_\lambda : W_j^\sigma \rightarrow W_j^\sigma$ and all its derivatives are smooth in λ at and near infinity.

Further, Ξ_λ extends to the double \widehat{W}_0^σ and the analogous properties hold.

Note that since W_j^σ is not an affine space, we do not have a natural notion of higher derivatives. However, we can think of W_j^σ as naturally being embedded as a submanifold of the linear space \widehat{W}_j , where we remove the conditions $s \geq 0$ and $\|\phi\|_{L^2} = 1$. (The same is true for the double \widehat{W}_j^σ .) Of course, $\widehat{W}_j \cong W_j \times \mathbb{R}$. We will then treat Ξ_λ as a map from W_j^σ to \widehat{W}_j . After defining Ξ_λ , it will be clear that this extends naturally to an L_j^2 neighborhood $\nu(W_j^\sigma)$ of W_j^σ in \widehat{W}_j . We can make sense of higher derivatives of Ξ_λ on $\nu(W_j^\sigma)$, since this is an open submanifold of a linear space, and these are the derivatives we will talk about (and measure in norm). Further, since we may consider elements of W_j^σ as elements of the linear space \widehat{W}_j , it makes sense to take differences of elements there and take their L_j^2 norms.

We will give the proof of Lemma 11.1 in Sections 11.1 and 11.2 below. We will then use this to obtain a number of important technical consequences for diffeomorphisms of path spaces. Note that given a path γ in W_j^σ , we may apply Ξ_λ slicewise to obtain a new path $\Xi_\lambda(\gamma)$ in W_j^σ . We will study the regularity of Ξ_λ on four-dimensional configurations. Before restating it, we need some discussion and terminology.

In parts (ii)-(iv) of the following proposition, we will discuss the smoothness of $\Xi_\lambda : W_j^\tau(x_\lambda, y_\lambda) \rightarrow W_j^\tau(x_\infty, y_\infty)$ in λ . This a priori is not defined, as the domain space is

changing. We postpone the discussion for what we mean by this in Section 11.3.3 (given after the statement of Proposition 11.13) for expediency.

Proposition 11.2. *Let Ξ_λ be as above. Fix stationary points x_∞ and y_∞ of $\mathcal{X}_q^{\text{gC},\sigma}$ in \mathcal{N} and let x_λ and y_λ be the stationary points of $\mathcal{X}_{q_\lambda}^{\text{gC},\sigma}$ which minimize L^2 distance to x_∞ and y_∞ respectively. Then for $\lambda \gg 0$ and $1 \leq j \leq k$, we have the following:*

- (i) *for a compact interval $I \subseteq \mathbb{R}$, the map Ξ_λ induces an S^1 -equivariant diffeomorphism of $\widetilde{W}_j^\tau(I \times Y)$ which is smooth in λ at and near ∞ and preserves $W_j^\tau(I \times Y)$,*
- (ii) *Ξ_λ induces diffeomorphisms from $\widetilde{W}_j^\tau(x_\lambda, y_\lambda)$ to $\widetilde{W}_j^\tau(x_\infty, y_\infty)$, which vary smoothly in λ at and near ∞ ,*
- (iii) *Ξ_λ induces a diffeomorphism from $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\lambda], [y_\lambda])$ to $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty])$,*
- (iv) *the diffeomorphisms $\Xi_\lambda : \widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\lambda], [y_\lambda]) \rightarrow \widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty])$ from the above item lift to smooth (in domain and also with respect to λ at and near ∞) bundle maps*

$$\begin{array}{ccc} \mathcal{V}_j^{\text{gC},\tau} & \xrightarrow{(\Xi_\lambda)_*} & \mathcal{V}_j^{\text{gC},\tau} \\ \downarrow & & \downarrow \\ \widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\lambda], [y_\lambda]) & \xrightarrow{\Xi_\lambda} & \widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty]). \end{array}$$

If $[x_\infty] \neq [y_\infty]$, the analogous statement also holds for $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$.

As for the case of W_j^σ , there exists a similar way to discuss the derivatives of Ξ_λ in four dimensions. We will think of $\widetilde{W}_j^\tau(I \times Y)$ as naturally being embedded as a submanifold of the linear space $\widehat{W}_j(I \times Y)$, where we remove the conditions $s(t) \geq 0$ and $\|\phi(t)\|_{L^2(Y)} = 1$. Following the discussion above, it turns out that Ξ_λ will extend to a neighborhood of $\widetilde{W}_j^\tau(I \times Y)$ in $\widehat{W}_j(I \times Y)$. For simplicity, we will work in the larger affine space to compute derivatives and measure the distance between configurations. This will not affect any statements about smoothness or regularity.

Before giving the proofs of Lemma 11.1 and Proposition 11.2, let us state two immediate corollaries of the latter. The first follows from part (i), by continuity:

Corollary 11.3. *For $1 \leq j \leq k$, if a sequence $\gamma_n \in W_{j,\text{loc}}^\tau(I \times Y)$ converges to some γ_∞ , then $\Xi_{\lambda_n}(\gamma_n) \rightarrow \gamma_\infty$ for any sequence $\lambda_n \rightarrow \infty$.*

The second corollary follows from part (ii) of Proposition 11.2:

Corollary 11.4. *For $1 \leq j \leq k$, $\gamma_0 \in W_j^\tau(x_\infty, y_\infty)$, if a sequence $\gamma_n \in W_j^\tau(x_{\lambda_n}, y_{\lambda_n})$ with $\lambda_n \rightarrow \infty$ satisfies*

$$\|\gamma_n - \Xi_{\lambda_n}^{-1}(\gamma_0)\|_{L_j^2(\mathbb{R} \times Y)} \rightarrow 0,$$

then

$$\|\Xi_{\lambda_n}(\gamma_n) - \gamma_0\|_{L_j^2(\mathbb{R} \times Y)} \rightarrow 0.$$

We now give the organization of the rest of this section. In Section 11.1, we give the construction of Ξ_λ . In Section 11.2, we prove the desired properties of Ξ_λ in Lemma 11.1. In Section 11.3, we prove parts (i) and (ii) of Proposition 11.2. Finally, in Section 11.4 we prove parts (iii) and (iv).

11.1. The construction of Ξ_λ .

Lemma 11.5. *For $\lambda \gg 0$, there exists an S^1 -equivariant diffeomorphism $\Xi_\lambda : W_0^\sigma \rightarrow W_0^\sigma$ which sends x_λ to x_∞ for each stationary point x_∞ of $\mathcal{X}_q^{\text{gC},\sigma}$ in \mathcal{N} .*

The construction will be similar to that of the diffeomorphism T_λ used in Section 8.2 to define F_λ . Since the blow-up is not a linear space, there is not a notion of translation like for the definition of T_λ , so we will work in charts (in the directions orthogonal to the S^1 orbits). The reader may note that here we are working with W_0^σ , while in the definitions of T_λ and F_λ from Section 8, we only worked with W_{k-1} . The explanation for this is that the function F_λ incorporated c , which is not well-behaved as a map from $W_0 \rightarrow W_0$ (since it contains quadratic terms), so we would not have been able to analyze F_λ in lower Sobolev regularity.

A more minor distinction is the following. To construct $T_\lambda(x)$, we used the function ω_λ defined in (181) to find which element of the S^1 -orbit of an approximate stationary point x_λ the point x is closest to, and then translate x by $\omega_\lambda(x)(x_\infty - x_\lambda)$, where x_∞ minimizes L^2 distance to x_λ . To define $\Xi_\lambda(x)$, we will compare x to x_∞ instead of x_λ , and then translate. The reason for this is that we will work with local charts, and it is simplest to work in charts around the same point (x_∞) for every λ , as opposed to defining Ξ_λ in terms of different charts for each λ and tracking the changes.

For notation, we choose representatives $x_\infty^1, \dots, x_\infty^m$ for the m orbits of stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$ in \mathcal{N} .

Proof of Lemma 11.5. We will first define Ξ_λ on an S^1 -invariant neighborhood of x_∞^i in W_0^σ for a fixed i ; for now, we omit the index i . Write $x_\infty = (a_\infty, s_\infty, \phi_\infty)$. Consider the submanifold

$$U_{x_\infty} = \{(a, s, \phi) \in W_0^\sigma \mid \text{Re}\langle \phi, i\phi_\infty \rangle_{L^2} = 0, \langle \phi, \phi_\infty \rangle_{L^2} > 0\}.$$

Note that U_{x_∞} is not open, as it only consists of configurations which are (real) orthogonal to the S^1 -orbit of x_∞ . We will define an S^1 -equivariant self-diffeomorphism of $S^1 \cdot U_{x_\infty}$ which takes x_λ to x_∞ and is the identity outside of a smaller S^1 -invariant neighborhood of x_∞ . We will then obtain the desired diffeomorphism Ξ_λ by repeating this construction for each orbit of stationary points of $\mathcal{X}_q^{\text{gC},\sigma}$, and taking Ξ_λ to be the identity outside of these neighborhoods.

First, we observe that there exists a diffeomorphism G_{x_∞} from U_{x_∞} to the Hilbert manifold with boundary

$$V_{x_\infty} = \{(a, s, \phi) \in (\ker d^*)_0 \oplus \mathbb{R} \oplus L^2(Y; \mathbb{S}) \mid s \geq 0, \langle \phi, \phi_\infty \rangle_{L^2} = 0\},$$

given by

$$G_{x_\infty} : (a, s, \phi) \mapsto \left(a, s, \frac{\phi}{\langle \phi, \phi_\infty \rangle_{L^2}} - \phi_\infty \right).$$

We remark that V_{x_∞} is a submanifold of $\mathcal{T}_{0,x_\infty}^{\text{gC},\sigma}$. We have that $G_{x_\infty}(x_\infty) = (a_\infty, s_\infty, 0)$. Denote this vector by z_∞ . Note also that in the definition of G_{x_∞} , have that $\text{Re}\langle \phi, \phi_\infty \rangle_{L^2} = \langle \phi, \phi_\infty \rangle_{L^2}$, since ϕ is orthogonal to the S^1 -orbit of ϕ_∞ . Observe also that the inverse of G_{x_∞} is given by the formula

$$G_{x_\infty}^{-1} : (a, s, \phi) \mapsto \left(a, s, \frac{\phi + \phi_\infty}{\|\phi + \phi_\infty\|_{L^2}} \right).$$

Since $x_\lambda = (a_\lambda, s_\lambda, \phi_\lambda)$ minimizes the L^2 distance from x_∞ to x_λ , it is necessarily orthogonal to the S^1 -orbit of x_∞ . Further, for $\lambda \gg 0$, $\|x_\infty - x_\lambda\|_{L^2}$ is arbitrarily small, and thus $\langle \phi_\lambda, i\phi_\infty \rangle_{L^2} > 0$, so x_λ is contained in U_{x_∞} . Let z_λ denote the image of x_λ under G_{x_∞} . We will construct an interpolation between translation by $z_\infty - z_\lambda$ (which may not be defined near the boundary of V_{x_∞} if x_∞ is irreducible) and the identity.

Pick $0 < \delta \ll \frac{1}{2}$ (independent of λ), such that an L^2 ball of size δ around the origin in V_{x_∞} has the following properties for $\lambda \gg 0$. First, this ball must contain $G_{x_\infty}(x_\lambda)$ and be disjoint from a δ -ball around $G_{x_\infty}(x_\infty^n)$ or $G_{x_\infty}(x_\lambda^n)$, should they be defined, for any $n \neq i$ in \mathcal{N} . Further, if x_∞ is irreducible, we choose δ such that the δ ball is contained in the interior of V_{x_∞} . By our choice of δ , if $(a, s, \phi) \in U_{x_\infty}$ satisfies $\langle \phi, \phi_\infty \rangle_{L^2} < \frac{1}{2}$, then $G_{x_\infty}(a, s, \phi)$ is not in this ball. It is in this ball where we will interpolate between the identity and the translation from x_λ to x_∞ .

Let $\alpha : [0, \infty) \rightarrow [0, 1]$ be a smooth bump function which is 1 on $[0, \delta/2]$ and 0 on $[\delta, \infty)$. We can now define a map Υ_λ from V_{x_∞} to V_{x_∞} given by

$$(240) \quad \Upsilon_\lambda : z \mapsto (1 - \alpha(\|\psi\|_{L^2}))z + \alpha(\|\psi\|_{L^2})(z + z_\infty - z_\lambda),$$

where we write $z = (b, r, \psi)$.

By the choice of δ , we have that this induces a well-defined map on V_{x_∞} , since in the case that x_∞ is reducible, the s -component is unchanged, and when x_∞ is irreducible, we have chosen the δ -ball such that translation does not happen near the boundary. Again, we point out that this map is translation by $z_\infty - z_\lambda$ inside of an L^2 ball of size $\delta/2$ in V_{x_∞} and is the identity outside of a δ ball. In particular, we see that z_λ is taken to z_∞ .

By the work of Section 7, we have that

$$(241) \quad x_\lambda \rightarrow x_\infty \text{ in } L_j^2 \text{ for all } j \text{ as } \lambda \rightarrow \infty$$

and thus

$$(242) \quad z_\lambda \rightarrow z_\infty \text{ in } L_j^2 \text{ for all } j \text{ as } \lambda \rightarrow \infty.$$

We claim that for $\lambda \gg 0$, the map Υ_λ in (240) induces a diffeomorphism of V_{x_∞} . Indeed, since $z_\lambda \rightarrow 0$, for $\lambda \gg 0$, the derivative of Υ_λ is close to the identity, and thus is an isomorphism at each point. The map is thus a local diffeomorphism, which is the identity outside of a ball. This is necessarily a self-covering map of a simply-connected space, and so we deduce that this is a global diffeomorphism.

We can now use G_{x_∞} to define the self-diffeomorphism Ξ_λ on U_{x_∞} . Since $G_{x_\infty}(x_\lambda) = z_\lambda$ and $G_{x_\infty}(x_\infty) = z_\infty$, we see that

$$\Xi_\lambda := G_{x_\infty}^{-1} \circ \Upsilon_\lambda \circ G_{x_\infty}$$

takes x_λ to x_∞ .

It follows by the construction that x_∞ is taken to x_λ . We then extend Ξ_λ to a diffeomorphism on the S^1 orbit of U_{x_∞} as follows. Define a function $\omega_\infty : S^1 \cdot U_{x_\infty} \rightarrow S^1$, similar to ω_λ^j in (181), by

$$(243) \quad \omega_\infty(x) = \frac{\operatorname{Re}\langle \phi, \phi_\infty \rangle_{L^2} + i \operatorname{Re}\langle \phi, i\phi_\infty \rangle_{L^2}}{((\operatorname{Re}\langle \phi, \phi_\infty \rangle_{L^2})^2 + (\operatorname{Re}\langle \phi, i\phi_\infty \rangle_{L^2})^2)^{1/2}}.$$

Note that ω_∞ has the property that if $x \in S^1 \cdot U_{x_\infty}$, then $\overline{\omega_\infty(x)} \cdot x \in U_{x_\infty}$. Therefore, we extend Ξ_λ to the S^1 orbit of U_{x_∞} by conjugating by ω_∞ :

$$x \mapsto \omega_\infty(x) \cdot \Xi_\lambda(\overline{\omega_\infty(x)} \cdot x).$$

By construction this extension is S^1 -equivariant. By repeating the above construction for each stationary point x_∞^i , we extend Ξ_λ to neighborhoods of every stationary point of \mathcal{N} . Note that Ξ_λ is the identity near the boundary of each such neighborhood. We finally extend Ξ_λ to a diffeomorphism of all of W_0^σ by the identity outside of these neighborhoods.

It is now clear from the construction that Ξ_λ is an S^1 -equivariant diffeomorphism of W_0^σ . \square

For future reference, we describe an explicit formula for Ξ_λ . Let us first introduce some notation to help express Ξ_λ more compactly. For each $i = 1, \dots, m$, define a function $\beta_i : W_0^\sigma \rightarrow [0, 1]$ by

$$\beta_i(x) = \begin{cases} 0 & \text{if } x = (a, s, \phi) \notin S^1 \cdot U^{x_\infty^i} \\ \alpha(\|\frac{\phi}{\langle \phi, \omega_\infty^i(x) \phi_\infty^i \rangle_{L^2}} - \omega_\infty^i(x) \phi_\infty^i\|_{L^2}) & \text{if } x = (a, s, \phi) \in S^1 \cdot U^{x_\infty^i}. \end{cases}$$

Further write $x_\lambda^i = (a_\lambda^i, s_\lambda^i, \phi_\lambda^i)$, where x_λ^i is the stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ corresponding to x_∞^i . We also write $v_\lambda^i = \frac{\phi_\lambda^i}{\langle \phi_\lambda^i, \phi_\infty^i \rangle_{L^2}} - \phi_\infty^i$. Note that $G_{x_\infty^i}(x_\lambda^i) = (a_\lambda^i, s_\lambda^i, v_\lambda^i)$. We get

$$(244) \quad \begin{aligned} \Xi_\lambda : W_0^\sigma &\rightarrow W_0^\sigma \\ (a, s, \phi) &\mapsto \left(a + \sum_i \beta_i(x)(a_\infty^i - a_\lambda^i), s + \sum_i \beta_i(x)(s_\infty^i - s_\lambda^i), \right. \\ &\quad \left. \frac{\phi - \sum_i \beta_i(x) \langle \phi, \phi_\infty^i \rangle_{L^2} \cdot v_\lambda^i}{\|\phi - \sum_i \beta_i(x) \langle \phi, \phi_\infty^i \rangle_{L^2} \cdot v_\lambda^i\|_{L^2}} \right). \end{aligned}$$

Here we take as notational convention that we do not worry about the well-definedness of ω_∞^i when β_i is zero. Observe that for each x , $\beta_i(x)$ is non-zero for at most one i .

We now discuss the extension of Ξ_λ to the double as claimed in the lemma. Recall that a stationary point x_∞ is reducible if and only if the approximate stationary point x_λ is. From (244), we therefore see that for $x \in W_0^\sigma$ with $s \ll 1$ (i.e. s much smaller than the smallest value of s_∞ for x_∞ irreducible), the map Ξ_λ preserves the s -component. In particular, Ξ_λ preserves and is tangent to the reducible locus. It follows that Ξ_λ extends to the double \widetilde{W}_0^σ as claimed. It will be clear that the arguments below establishing the desired properties of Ξ_λ for W_0^σ extend to \widetilde{W}_0^σ .

Remark 11.6. It is worth noting that almost all of the pieces of Ξ_λ are just determined by the spinorial component of x . Further, in the end result, the formula for Ξ_λ only uses the function ω in the bump-like functions β_i .

11.2. Three-dimensional properties of Ξ_λ . We now show that Ξ_λ has the stated properties in Lemma 11.1. After that, we will establish some additional bounds that will be useful for Proposition 11.2.

We begin with the first item in Lemma 11.1.

Proof of Lemma 11.1(i). It is clear from the construction of Ξ_λ , that each x_λ^i is mapped to x_∞^i . The S^1 -equivariance then implies that x_λ is mapped to x_∞ for any stationary point x_λ of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ in \mathcal{N} . \square

11.2.1. Smoothness properties.

Proof of Lemma 11.1(ii). First, note that the functions β_i , ω_∞^i , and $\phi \mapsto \langle \phi, \phi_\infty^i \rangle_{L^2}$ are smooth on L_j^2 for any $1 \leq j \leq k$, since the L^2 inner-product is, and $x_\infty^i \in L_k^2$. At each step

of the construction of Ξ_λ , where Ξ_λ is not the identity (or in other words, some β_i is non-zero), we are simply scalar multiplying by ω_∞^i , which is non-zero, or by adding a multiple of β_i or $\langle \phi, \phi_\infty^i \rangle_{L^2}$ times x_∞^i or ϕ_∞^i . Since x_λ^i, x_∞^i are in L_k^2 , these (invertible) operations necessarily preserve the condition of being an L_j^2 configuration for $1 \leq j \leq k$, and we have the desired result. \square

At the beginning of this section, Ξ_λ was claimed to extend to a neighborhood of W_j^σ in \widehat{W}_j . This is clear from the explicit description of Ξ_λ given in (244) since wherever Ξ_λ disagrees with the identity, we have that $\beta_i \neq 0$ for precisely one i , and in this case $\langle \phi, \omega_\infty^i \phi_\infty^i \rangle_{L^2}$ is real and positive (in fact, at least $1/2$); from this it is easy to see that Ξ_λ extends to a neighborhood of this point in \widehat{W}_j , when we also expand the target to \widehat{W}_j as well.

The following completes the proof of Lemma 11.1.

Proof of Lemma 11.1(iii). Since Ξ_λ is defined by extending a self-diffeomorphism of $U_{x_\infty^i}$ by S^1 , and the space $U_{x_\infty^i}$ is defined independently of λ , it suffices to show that $\Xi_{f^{-1}(r)}|_{U_{x_\infty^i}}$ is smooth at and near $r = 0$. Recall from Corollary 7.8 that the correspondence $r \mapsto [x_{f^{-1}(r)}]$ is differentiable at and near $r = 0$. It is easy to check that this lifts to a smoothness statement without quotienting by S^1 , due to the condition that $\text{Re}\langle x_{f^{-1}(r)}, ix_\infty^i \rangle_{L^2} = 0$. Using (244), we see the only terms in $\Xi_{f^{-1}(r)}$ or its derivatives which depend on r are smooth functions of $x_{f^{-1}(r)}$, and thus differentiable at and near $r = 0$. This gives the desired smoothness statement. \square

11.2.2. Some three-dimensional bounds on Ξ_λ . Before moving on to the four-dimensional properties of Ξ_λ in Proposition 11.2, we study a few additional properties of Ξ_λ . For notational simplicity, we work in a neighborhood $S^1 \cdot U_{x_\infty^i}$ for some i . For the rest of this subsection, we again omit the index i from the notation.

It is clear from (244) that the complication in Ξ_λ is in the spinorial component. We begin by simplifying the notation and then establishing a few key bounds. We write

$$(245) \quad f_\lambda(x) = \phi - \langle \phi, \phi_\infty \rangle_{L^2} \cdot \beta(x) \cdot v_\lambda$$

and thus we can express the spinorial component of Ξ_λ as $\frac{f_\lambda}{\|f_\lambda\|_{L^2}}$ for $x \in S^1 \cdot U_{x_\infty}$. Therefore, Ξ_λ can be written a bit more succinctly as

$$(246) \quad \Xi_\lambda(a, s, \phi) = \left(a + \beta(x)(a_\infty - a_\lambda), \ s + \beta(x)(s_\infty - s_\lambda), \ \frac{f_\lambda(x)}{\|f_\lambda(x)\|_{L^2}} \right).$$

We point out some bounds that will be useful for us when proving Proposition 11.2. By the work of Section 7, we have that $\|x_\lambda - x_\infty\|_{L_k^2} \rightarrow 0$. It is then easy to see that for $\lambda \gg 0$, since $\|\phi_\infty\|_{L^2} = 1$, we have

$$(247) \quad \|v_\lambda\|_{L^2} \rightarrow 0.$$

From (247), using that the spinorial component of an element of W^σ has unit L^2 -norm, we can also deduce that for $\lambda \gg 0$:

$$(248) \quad \frac{1}{2} \leq \|f_\lambda(x)\|_{L^2} \leq \frac{3}{2}.$$

11.3. Four-dimensional properties of Ξ_λ . In this section, we prove Proposition 11.2. For the rest of this section, we let j be an integer such that $1 \leq j \leq k$. Before proceeding, we remind the reader that the L_j^2 -norm of a four-dimensional configuration γ on $I \times Y$ is expressed as

$$\|\gamma\|_{L_j^2(I \times Y)}^2 = \sum_{n=0}^j \int_a^b \left\| \frac{d^n}{dt^n} \gamma(t) \right\|_{L_{j-n}^2(Y)}^2 dt.$$

11.3.1. Sobolev multiplication and superposition operators. In order to study how Ξ_λ acts on paths in $W_j^T(I \times Y)$, we will need some elementary variants of Sobolev multiplication and superposition operators. Many of these are well-known or easily deduced from the definitions and Sobolev multiplication. We include the proofs so as to reference them for analogous results later on that are less standard.

Lemma 11.7. *Let $I \subseteq \mathbb{R}$. Then, taking inner products induces smooth maps*

$$(249) \quad \widehat{W}_j(I \times Y) \times \widehat{W}_j(I \times Y) \rightarrow L_j^2(I; \mathbb{R}), \quad (\gamma(t), \eta(t)) \mapsto \langle \gamma(t), \eta(t) \rangle_{L^2(Y)},$$

$$(250) \quad \widehat{W}_j(I \times Y) \times \widehat{W}_j \rightarrow L_j^2(I; \mathbb{R}), \quad (\gamma, x) \mapsto \langle \gamma, x \rangle_{L^2(Y)}.$$

Proof. While the smoothness of (249) is a special case of Sobolev multiplication, we provide a proof, since it makes the second claim easier to justify and we will use similar techniques throughout the rest of the section. We compute

$$\begin{aligned} \|\langle \gamma, \eta \rangle_{L^2(Y)}\|_{L_j^2(I)}^2 &= \sum_{n=0}^j \left\| \frac{d^n}{dt^n} \langle \gamma, \eta \rangle_{L^2(Y)} \right\|_{L^2(I)}^2 \\ &= \sum_{n=0}^j \int_I \left| \sum_{i=0}^n \binom{n}{i} \left\langle \frac{d^i}{dt^i} \gamma, \frac{d^{n-i}}{dt^{n-i}} \eta \right\rangle_{L^2(Y)} \right|^2 dt \\ &\leq \sum_{n=0}^j \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} \int_I \left| \left\langle \frac{d^{i_1}}{dt^{i_1}} \gamma, \frac{d^{n-i_1}}{dt^{n-i_1}} \eta \right\rangle_{L^2(Y)} \cdot \left\langle \frac{d^{i_2}}{dt^{i_2}} \gamma, \frac{d^{n-i_2}}{dt^{n-i_2}} \eta \right\rangle_{L^2(Y)} \right| dt \\ &\leq \sum_{n=0}^j \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} \int_I \left\| \frac{d^{i_1}}{dt^{i_1}} \gamma \right\|_{L^2(Y)} \left\| \frac{d^{n-i_1}}{dt^{n-i_1}} \eta \right\|_{L^2(Y)} \left\| \frac{d^{i_2}}{dt^{i_2}} \gamma \right\|_{L^2(Y)} \left\| \frac{d^{n-i_2}}{dt^{n-i_2}} \eta \right\|_{L^2(Y)} dt. \end{aligned}$$

Let us focus on a fixed term in the final term above, i.e., we fix a value of n, i_1, i_2 . Note that $n \leq j$, so one of $n - i_1$ or i_1 is bounded above by $\frac{j}{2}$. Without loss of generality, it is i_1 . Since $j \geq 1$, we have that

$$j - i_1 \geq \frac{j}{2} \geq \frac{1}{2}.$$

(In fact, $j - i_1 \geq 1$, because j and i_1 are integers.) We deduce that the composition of $\frac{d^{i_1}}{dt^{i_1}} : \widehat{W}_j(I \times Y) \rightarrow \widehat{W}_{j-i_1}(I \times Y)$ with the restriction map to a slice $\widehat{W}_{j-i_1}(I \times Y) \rightarrow L^2(\{t\} \times Y)$ is continuous. From here we see that there exists a constant C_{i_1} , depending only on Y and I , such that

$$\left\| \frac{d^{i_1}}{dt^{i_1}} \gamma(t) \right\|_{L^2(Y)} \leq C_{i_1} \|\gamma\|_{L_j^2(I \times Y)}, \quad \text{for all } t \in I.$$

We can repeat the same argument for i_2 or $n - i_2$. Again, without loss of generality we assume that $i_2 \leq \frac{j}{2}$, and we obtain bounds

$$\left\| \frac{d^{i_2}}{dt^{i_2}} \gamma(t) \right\|_{L^2(Y)} \leq C_{i_2} \|\gamma\|_{L_j^2(I \times Y)}, \quad \text{for all } t \in I.$$

Therefore, we may write

$$\begin{aligned} \|\langle \gamma, \eta \rangle_{L^2(Y)}\|_{L_j^2(I)}^2 &\leq \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} \int_I \left\| \frac{d^{i_1}}{dt^{i_1}} \gamma \right\|_{L^2(Y)} \left\| \frac{d^{n-i_1}}{dt^{n-i_1}} \eta \right\|_{L^2(Y)} \left\| \frac{d^{i_2}}{dt^{i_2}} \gamma \right\|_{L^2(Y)} \left\| \frac{d^{n-i_2}}{dt^{n-i_2}} \eta \right\|_{L^2(Y)} dt \\ &\leq \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} C_{i_1} C_{i_2} \|\gamma\|_{L_j^2(I \times Y)}^2 \int_I \left\| \frac{d^{n-i_1}}{dt^{n-i_1}} \eta \right\|_{L^2(Y)} \left\| \frac{d^{n-i_2}}{dt^{n-i_2}} \eta \right\|_{L^2(Y)} dt \\ &\leq \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} C_{i_1} C_{i_2} \|\gamma\|_{L_j^2(I \times Y)}^2 \left\| \frac{d^{n-i_1}}{dt^{n-i_1}} \eta \right\|_{L^2(I \times Y)} \left\| \frac{d^{n-i_2}}{dt^{n-i_2}} \eta \right\|_{L^2(I \times Y)} \\ &\leq \sum_{i_1, i_2=0}^n \binom{n}{i_1} \binom{n}{i_2} C_{i_1} C_{i_2} \|\gamma\|_{L_j^2(I \times Y)}^2 \|\eta\|_{L_j^2(I \times Y)}^2, \end{aligned}$$

where in the penultimate inequality we have used Cauchy-Schwarz. An analogous inequality holds if we instead replace i_1 or i_2 with $n - i_1$ or $n - i_2$ respectively, depending on their values.

By summing over the relevant constants, we obtain a bound

$$\|\langle \gamma, \eta \rangle_{L^2(Y)}\|_{L_j^2(I)} \leq C \|\gamma\|_{L_j^2(I \times Y)} \|\eta\|_{L_j^2(I \times Y)},$$

where C is independent of γ, η . It follows that slicewise inner products define a *bounded*, bilinear map. Such a map is necessarily smooth.

For (250), we can repeat the proof above. Note that an element of \widehat{W}_j does not determine an element of $\widehat{W}_j(I \times Y)$ if I is not compact. However, for each term analyzed in the above argument, we only needed that one of the terms involved (either a derivative of γ or of η) be $L^2(Y)$ bounded uniformly in time, with the other term square-integrable in time. In this case, the three-dimensional configuration in \widehat{W}_j is constant in time, and has bounded $L^2(Y)$ norm (since $j \geq 0$), so we have the desired result. \square

Lemma 11.8. *Let $I \subseteq \mathbb{R}$. Then, pointwise multiplication induces smooth maps*

$$(251) \quad L_j^2(I; \mathbb{R}) \times \widehat{W}_j(I \times Y) \rightarrow \widehat{W}_j(I \times Y), \quad (g, \gamma) \mapsto g\gamma,$$

$$(252) \quad L_j^2(I; \mathbb{R}) \times \widehat{W}_j \rightarrow \widehat{W}_j(I \times Y), \quad (g, \gamma) \mapsto g\gamma.$$

Proof. The argument is similar to that of the above lemma. We begin with the first equation. Recall that for any $\gamma \in \widehat{W}_j(I \times Y)$, we have a norm for $\gamma(t, y)$, defined for any $(t, y) \in I \times Y$; we write this function as $|\gamma|$. We compute

$$\begin{aligned} \|g\gamma\|_{L_j^2(I \times Y)}^2 &= \sum_{n=0}^j \int_{I \times Y} \left| \sum_{i=0}^n \frac{d^i}{dt^i} (g) \nabla^{n-i} \gamma \right|^2 \\ &\leq \sum_{i_1, i_2=0}^n \binom{n}{i_1} \cdot \binom{n}{i_2} \int_{I \times Y} \left| \frac{d^{i_1}}{dt^{i_1}} (g) \nabla^{n-i_1} \gamma \right| \cdot \left| \frac{d^{i_2}}{dt^{i_2}} (g) \nabla^{n-i_2} \gamma \right|. \end{aligned}$$

We can now apply the same argument as in the first part of Lemma 11.7 to obtain a bound

$$\|g\gamma\|_{L_j^2(I \times Y)} \leq K \|g\|_{L_j^2(I)} \|\gamma\|_{L_j^2(I \times Y)},$$

where K is a constant independent of g and γ . This shows that pointwise multiplication is continuous. Smoothness again follows by bilinearity.

The second equation can be obtained by applying the same modifications to the proof as in Lemma 11.7 to establish (250). \square

The following lemma is standard. We include a proof for completeness.

Lemma 11.9. *Let $I \subseteq \mathbb{R}$ be an interval, and $h : \mathbb{R} \rightarrow \mathbb{R}$ a smooth function. Consider the transformation of $C^\infty(\mathbb{R})$ given by*

$$H : g \mapsto h \circ g.$$

(a) *If I is compact, then H induces a smooth map from $L_j^2(I; \mathbb{R})$ to $L_j^2(I; \mathbb{R})$.*

(b) *The same holds true for arbitrary I (not necessarily compact), provided that $h(0) = 0$.*

Proof. (a) We use Faà di Bruno's formula for the higher derivatives of a composition:

$$(253) \quad \frac{d^n}{dt^n} h(g(t)) = \sum C(m_1, \dots, m_n) \cdot h^{(m_1 + \dots + m_n)}(g(t)) \cdot \prod_{i=1}^n (g^{(i)}(t))^{m_i}.$$

Here, the sum is taken over all n -tuples of nonnegative integers (m_1, \dots, m_n) with

$$1 \cdot m_1 + 2 \cdot m_2 + \dots + n \cdot m_n = n,$$

and the constants $C(m_1, \dots, m_n)$ are

$$C(m_1, \dots, m_n) = \frac{n!}{m_1! 1!^{m_1} m_2! 2!^{m_2} \dots m_n! n!^{m_n}}.$$

Let $g \in L_j^2(I; \mathbb{R})$. We first want to show that $h \circ g$ is also in $L_j^2(I; \mathbb{R})$, that is, the n th derivative $\frac{d^n}{dt^n} h(g(t))$ is in $L^2(I; \mathbb{R})$, for all $0 \leq n \leq j$.

We claim that each summand in the expression (253) is in L^2 . Indeed, since g is in L_j^2 and $j \geq 1$, we have that g is continuous. Further, since h is smooth, we see that the expression $h^{(m_1 + \dots + m_n)}(g(t))$ is continuous in t , and therefore bounded (because I is compact). Moreover, unless $i = n$, the factors $g^{(i)}(t)$ are in L_1^2 ; hence, by Sobolev multiplication, their product is also in L_1^2 , and thus in L^2 . It follows that

$$(254) \quad h^{(m_1 + \dots + m_n)}(g(t)) \cdot \prod_{i=1}^n (g^{(i)}(t))^{m_i} \in L^2,$$

as long as $m_n = 0$, since $h^{(m_1 + \dots + m_n)} \circ g$ is bounded. In the special case $m_n > 0$, we see that $m_n = 1$ and all the other m_i must be zero, so $(g^{(i)}(t))^{m_i} = g^{(n)}(t)$ is still in L^2 , and we get the same conclusion.

We have shown that H maps $L_j^2(I; \mathbb{R})$ to $L_j^2(I; \mathbb{R})$. To see that H is smooth, we compute its derivatives:

$$(255) \quad (\mathcal{D}^r H)_g(\xi_1, \dots, \xi_r) = (h^{(r)} \circ g) \cdot \xi_1 \cdot \dots \cdot \xi_r.$$

When $g, \xi_1, \dots, \xi_r \in L_j^2(I; \mathbb{R})$, by applying the fact we just proved with $h^{(r)}$ instead of h , and making use of the Sobolev multiplication $L_j^2 \times L_j^2 \rightarrow L_j^2$, we get that $(\mathcal{D}^r H)_g(\xi_1, \dots, \xi_r)$ is in L_j^2 , as desired.

(b) If I is non-compact, we still have that $h^{(m_1+\dots+m_n)}$ and g are continuous. Further, by the Sobolev embedding theorem, $g \in L_j^2$ implies that g is bounded. Therefore, it is still the case that $h^{(m_1+\dots+m_n)} \circ g$ is bounded. We also have the Sobolev multiplication $L_j^2 \times L_j^2 \rightarrow L_j^2$ for $j \geq 1$, so the same arguments as before show that (254) holds, provided that not all m_i are zero.

When $n = 0$ and all m_i are zero, the constant function $1 = \prod_{i=1}^n (g^{(i)}(t))^{m_i}$ is not in L^2 . Nevertheless, we claim that $h^{(m_1+\dots+m_n)}(g(t)) \cdot 1 = h(g(t))$ is in L^2 . To see this, we make use of the hypothesis $h(0) = 0$. The derivative of h at 0 is

$$(256) \quad h'(0) = \lim_{t \rightarrow 0} \frac{h(t)}{t} \in \mathbb{R}.$$

Recall that g is bounded, so there is some $K > 0$ such that $|g(t)| \leq K$. Let M be the supremum of $h(t)/t$ over the interval $[-K, K]$. From (256) and the continuity of h , we see that this supremum is finite. Therefore, we have

$$|h(g(t))| \leq M|g(t)|,$$

for all $t \in \mathbb{R}$. Since $g \in L^2$, this easily implies that $h \circ g$ is in L^2 .

We have now shown that H maps $L_j^2(I; \mathbb{R})$ to $L_j^2(I; \mathbb{R})$. To check that H is smooth, it remains to show that the expression (255) is in $L_j^2(I; \mathbb{R})$, when $g, \xi_1, \dots, \xi_r \in L_j^2(I; \mathbb{R})$.

By Sobolev multiplication, the product $\xi = \xi_1 \dots \xi_r$ is in L_j^2 . For $0 \leq n \leq j$, we need to verify that

$$(257) \quad \frac{d^n}{dt^n} ((h^{(r)} \circ g) \cdot \xi) = \sum_{s=0}^n \binom{n}{s} \frac{d^s}{dt^s} (h^{(r)} \circ g) \cdot \frac{d^{n-s}}{dt^{n-s}} \xi$$

is in L^2 .

For $s \neq 0$, the arguments above (based on Faà di Bruno's formula), applied to $h^{(r)}$ instead of h , show that the s^{th} derivative of $h^{(r)} \circ g$ is in L^2 . Further, since $n - s < n \leq j$, the $(n - s)^{\text{th}}$ derivative of ξ is in L_1^2 . It follows that (257) is in L^2 .

In the special case $s = 0$, the expression $h^{(r)} \circ g$ is bounded, and $\frac{d^n}{dt^n} \xi \in L^2$, so we again get that their product is in L^2 . This concludes the proof. \square

Let $I \subseteq \mathbb{R}$ and $\epsilon > 0$. Consider the subset $L_j^2(I; \mathbb{R}_{>\epsilon})$ inside $L_j^2(I; \mathbb{R})$ consisting of functions with values in (ϵ, ∞) . Note that for this subset to be non-empty, the interval I must be compact. Further, since $j \geq 1$, $L_j^2(I; \mathbb{R}_{>\epsilon})$ is an open subset of $L_j^2(I; \mathbb{R})$.

Lemma 11.10. *Let $I \subseteq \mathbb{R}$. The map $x \mapsto \frac{1}{x}$ induces a smooth map $H : L_j^2(I; \mathbb{R}_{>\epsilon}) \rightarrow L_j^2(I; \mathbb{R})$, given by $H(g)(x) = \frac{1}{g(x)}$. A similar statement applies to $x \mapsto \sqrt{x}$.*

Proof. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be any smooth function such that $h(x) = 1/x$ for $x \geq \epsilon$. Then, the result follows from Lemma 11.9(a), applied to h . The same argument works for \sqrt{x} . \square

With this, we can study the regularity of the L^2 norm as well.

Lemma 11.11. *Let $I \subseteq \mathbb{R}$. Let $\gamma \in \widehat{W}_j(I \times Y)$ have $\|\gamma(t)\|_{L^2(Y)} \geq \epsilon > 0$ for all t . For a neighborhood of γ in $\widehat{W}_j(I \times Y)$, the association $\gamma \mapsto \|\gamma\|_{L^2(Y)}$ is a smooth map to $L_j^2(I; \mathbb{R})$.*

Proof. Writing $\|\gamma\|_{L^2(Y)} = \langle \gamma, \gamma \rangle_{L^2(Y)}^{\frac{1}{2}}$, the result follows from Lemma 11.7 and the second statement in Lemma 11.10. \square

Finally, we have

Lemma 11.12. *Let $I \subset \mathbb{R}$ be compact. Postcomposition with the bump function α induces a smooth map $L_j^2(I; \mathbb{R}_{\geq 0}) \rightarrow L_j^2(I; \mathbb{R}_{\geq 0})$ given by $g \mapsto \alpha \circ g$.*

Proof. This is an immediate consequence of Lemma 11.9(a). \square

11.3.2. On L_j^2 regularity of Ξ_λ for compact cylinders. With the above technical lemmas established, we can now prove the first part of Proposition 11.2, which concerns the L_j^2 regularity of Ξ_λ on compact cylinders. First, recall from Lemma 11.1 that Ξ_λ extends to \widetilde{W}_j^σ . There is an obvious involution ι on \widetilde{W}_j^σ given by $(a, s, \phi) \mapsto (a, -s, \phi)$, such that $\Xi_\lambda(a, -s, \phi) = \Xi_\lambda(\iota(a, s, \phi)) = \iota(\Xi_\lambda(a, s, \phi))$. When discussing the terms involved in Ξ_λ , we will suppress this from the notation; for example, if (a, s, ϕ) has $s < 0$, we will write $\beta_i(a, s, \phi)$ instead of $\beta_i(\iota(a, s, \phi))$. Since Ξ_λ is ι -equivariant, it is easy to verify that this will not affect any smoothness or regularity.

Proof of Proposition 11.2(i). Let I be a compact interval. The idea is to apply the technical lemmas proved above on Sobolev multiplication and superposition operators to the explicit formula for Ξ_λ in (244). Fix $\gamma \in \widetilde{W}_j^T(I \times Y)$, where we write $\gamma(t) = (a(t), s(t), \phi(t))$. Recall that $\|\gamma(t)\|_{L^2(Y)} = 1$ for all $t \in I$. We break Ξ_λ into elementary pieces and study their regularity.

First, observe that $\langle \phi(t), \phi_\infty^i \rangle_{L^2(Y)} \in L_j^2(I; \mathbb{R})$ by Lemma 11.7. By Lemma 11.8, we have that $\langle \phi(t), \phi_\infty^i \rangle_{L^2(Y)} \cdot v_\lambda^i$ is in L_j^2 .

Note that whenever $\gamma(t) \in S^1 \cdot U_{x_\infty^i}$, we have $|\langle \phi(t), \phi_\infty^i \rangle_{L^2(Y)}| \geq \frac{1}{2}$ by our choice of the δ -balls in the definition of Ξ_λ . (Here we are using that there exists $u(t) \in S^1$ such that $u(t) \cdot \phi(t) \in U_{x_\infty^i}$.) Using Lemma 11.10, formula (243), and the Sobolev multiplication $L_j^2(I; \mathbb{R}) \times L_j^2(I; \mathbb{R}) \rightarrow L_j^2(I; \mathbb{R})$, we can deduce that $\omega_\infty^i(\gamma) \in L_j^2(I; \mathbb{C})$. Consequently, it follows from Lemmas 11.8, 11.10, and 11.12 that $\beta_i(\gamma) \in L_j^2(I; \mathbb{R})$. Therefore, we see by Lemma 11.8 that

$$(a + \beta_i(\gamma)(a_\infty^i - a_\lambda^i), s + \beta_i(\gamma)(s_\infty^i - s_\lambda^i), f_\lambda^i(\gamma))$$

is in $\widehat{W}_j(I \times Y)$. The slicewise L^2 norm of the spinorial component of this path is bounded below by $1/2$ by (248). Therefore, we see that

$$\left(a + \beta_i(\gamma)(a_\infty^i - a_\lambda^i), s + \beta_i(\gamma)(s_\infty^i - s_\lambda^i), \frac{f_\lambda^i(\gamma)}{\|f_\lambda^i(\gamma)\|_{L^2(Y)}} \right),$$

is contained in $\widehat{W}_j(I \times Y)$ by Lemmas 11.8, 11.10, and 11.11. This expression is $\Xi_\lambda(\gamma)$ by (246), as $\beta_n(\gamma) = 0$ for any $n \neq i$, when $\beta_i(\gamma) \neq 0$.

The smoothness of Ξ_λ follows from the relevant smoothness established in the sequence of technical lemmas above about Sobolev multiplication and superposition operators. To see that Ξ_λ is a diffeomorphism, we can apply similar arguments to show that Ξ_λ^{-1} satisfies the same regularity and smoothness properties.

We now show that Ξ_λ depends smoothly on λ , at and near infinity. If we change the definition of Ξ_λ by varying x_λ^i smoothly, the arguments above (i.e., the application of Lemmas 11.7-11.12) show that we will vary the induced diffeomorphism in a smooth way. This only applies for small variations of x_λ^i whose spinorial component stays slicewise in the L^2 ball of size δ around ϕ_∞^i , but that is all we will need. By Corollary 7.8, we see that

$x_{f^{-1}(r)}$ is smooth at and near $r = 0$, and thus we obtain that $\Xi_{f^{-1}(r)}$ is smooth at and near $r = 0$.

It remains to verify that Ξ_λ preserves $W_j^\tau(I \times Y)$ (i.e. Ξ_λ preserves the condition $s \geq 0$). This is trivial since Ξ_λ takes W_0^σ to W_0^σ . \square

11.3.3. On L_j^2 regularity of Ξ_λ . Before studying further regularity properties of Ξ_λ , we need to introduce some further notation, generalizing the construction of $W_j^\tau(x_\lambda, y_\lambda)$. Let $I \subseteq \mathbb{R}$ be an interval. Given a path $\zeta \in \widehat{W}_{j,loc}^\tau(I \times Y)$ (analogous to the smooth reference path γ_0 in defining $W_j^\tau(x_\lambda, y_\lambda)$), we define

$$\widehat{W}_j(I \times Y, \zeta) = \{\gamma \in \widehat{W}_{j,loc}^\tau(I \times Y) \mid \gamma - \zeta \in \widehat{W}_j(I \times Y)\}.$$

We equip this space with the L_j^2 metric (not $L_{j,loc}^2$). Note that many different functions can induce the same space. Further, note that $\widehat{W}_j(I \times Y, \zeta)$ is a Banach manifold, since we have an affine identification with $\widehat{W}_j(I \times Y)$.

Towards studying the map Ξ_λ on $\widehat{W}_j^\tau(x_\lambda, y_\lambda)$, we will begin with a regularity result for paths in \widehat{W}_0^σ that are contained entirely in some neighborhood of an orbit of stationary points (or their involutes under ι). As usual, for an element $x \in \widehat{W}_j$, we will abusively write x for the induced constant path in $\widehat{W}_{j,loc}(I \times Y)$.

Proposition 11.13. *Fix a stationary point x_∞^i of $\mathcal{X}_q^{\text{agC},\sigma}$ as in the construction of Ξ_λ , and let x_λ^i be the corresponding stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$. Let $T > 0$. Let $\gamma \in \widehat{W}_j^\tau([T, \infty) \times Y, x_\lambda^i)$ be such that $\beta_i(\gamma) \equiv 1$. Then,*

- (a) *in a neighborhood of γ in $\widehat{W}_j([T, \infty) \times Y, x_\lambda^i)$, Ξ_λ induces a smooth map to $\widehat{W}_j([T, \infty) \times Y, x_\infty^i)$, and*
- (b) *this family of maps is smooth in λ at and near infinity.*

An analogous result applies for the half-cylinder $(-\infty, T] \times Y$.

We postpone the proof of Proposition 11.13, but do explain what we mean by smoothness of Ξ_λ in λ , since the domain of the function is changing. Fix a reference path γ_0 in $W_{j,loc}^\tau([T, \infty) \times Y)$ which agrees with x_∞ for $t \gg 0$. Note that $\Xi_\lambda^{-1}(\gamma_0)$ provides an $L_{j,loc}^2$ reference path which agrees with x_λ for $t \gg 0$ by the work of Section 11.3.2. Then, Ξ_λ induces a map,

$$\widehat{W}_j([T, \infty) \times Y) \rightarrow \widehat{W}_j([T, \infty) \times Y), \quad \gamma - \Xi_\lambda^{-1}(\gamma_0) \mapsto \Xi_\lambda(\gamma) - \gamma_0,$$

defined in a neighborhood of $\gamma - \Xi_\lambda^{-1}(\gamma_0)$, where $\gamma \in \widehat{W}_{j,loc}^\tau([T, \infty) \times Y)$ satisfies $\gamma - \Xi_\lambda^{-1}(\gamma_0) \in \widehat{W}_j([T, \infty) \times Y)$. Since the domain and target of this map are constant in λ , we can make sense of smoothness in λ at and near ∞ . There is an analogous notion of smoothness in λ for maps from $\widehat{W}_j^\tau(x_\lambda, y_\lambda)$ to $\widehat{W}_j^\tau(x_\infty, y_\infty)$ (assuming they extend to neighborhoods in the relevant larger affine space), which is what is meant in Proposition 11.2(ii) and Proposition 11.13. More generally, given $f : \widehat{W}_{j,loc}(I \times Y) \rightarrow \widehat{W}_{j,loc}(I \times Y)$, one can analogously define smoothness in γ of the family of maps $f : \widehat{W}_j(I \times Y, \gamma) \rightarrow \widehat{W}_j(I \times Y, f(\gamma))$. One can check that these notions are independent of the choice of reference path.

Before proving Proposition 11.13, we use it to complete the proof of Proposition 11.2(ii).

Proof of Proposition 11.2(ii). Let $\gamma \in \widetilde{W}_j^\tau(x_\infty, y_\infty)$. First, we establish that $\Xi_\lambda(\gamma) \in \widetilde{W}_j^\tau(x_\infty, y_\infty)$. Fix a reference path γ_0 from x_∞ to y_∞ which is constant outside of $[-T, T]$. Write x_∞^i (respectively x_∞^n) for the indexed stationary point from the construction of Ξ_λ that is in the orbit of x_∞ (respectively y_∞).

We have that $\Xi_\lambda(\gamma) \in \widetilde{W}_{j,loc}^\tau(\mathbb{R} \times Y)$ by Proposition 11.2(i). Since $\gamma \in \widetilde{W}_j^\tau(x_\infty, y_\infty)$, we have that $\gamma - \gamma_0$ has finite $L_j^2(\mathbb{R} \times Y)$ norm, with $j \geq 1$. Thus, we see that for $t \gg 0$, we must have that $\|\gamma(t) - y_\infty\|_{L^2(Y)}$ is sufficiently small so that $\gamma(t)$ is contained in $S^1 \cdot U_{x_\infty^n}$ and $\beta_n(\gamma(t)) = 1$. A similar statement applies for $t \ll 0$, where we see that $\beta_i(\gamma(t)) = 1$.

By Proposition 11.13, we see that

$$\Xi_\lambda(\gamma|_{[T,\infty)}) \in \widehat{W}_j([T,\infty) \times Y, \Xi_\lambda(\gamma_0|_{[T,\infty)})).$$

A similar statement applies for $(-\infty, T] \times Y$ as well. Since $\gamma|_{[-T,T]} \in \widetilde{W}_{j,loc}^\tau([-T,T] \times Y)$, we can put these three pieces of γ together to see that $\gamma \in \widetilde{W}_j^\tau(x_\infty, y_\infty)$.

We will use a similar argument for smoothness. We provide the argument for the existence of the first derivative of Ξ_λ ; the higher derivatives are similar. Let $\eta \in T_\gamma \widehat{W}_j(\mathbb{R} \times Y, \gamma_0) \cong \widehat{W}_j(\mathbb{R} \times Y)$. We would like to see that

$$(258) \quad \lim_{h \rightarrow 0} \left\| \frac{\Xi_\lambda(\gamma + h \cdot \eta) - \Xi_\lambda(\gamma)}{h} \right\|_{L_j^2(\mathbb{R} \times Y)}$$

exists. This limit needs to be taken with respect to L_j^2 , and not $L_{j,loc}^2$. As before, we have that $\|\eta(t)\|_{L^2(Y)}$ is uniformly bounded in t , because of the $L_j^2(\mathbb{R} \times Y)$ bounds on η . Therefore, for h sufficiently small, we have that $\gamma + h \cdot \eta(t)$ is in $S^1 \cdot U_{x_\infty^n}$ for $t \geq T$, where we can choose T independent of h . A similar statement applies for $t \leq -T$. Therefore, by Proposition 11.13 the above limit exists in the L_j^2 topology when we replace $\mathbb{R} \times Y$ with $[T, \infty) \times Y$ and with $(-\infty, T] \times Y$. By Proposition 11.2(i), we have that the limit in (258) exists when we replace \mathbb{R} with $[-T, T]$. This establishes the existence of the limit on $L_j^2(\mathbb{R} \times Y)$.

The smoothness of Ξ_λ in λ at and near infinity can again be deduced by a similar argument, applying Proposition 11.2(i) for a fixed compact interval and Proposition 11.13 outside of this interval. \square

Before moving on to Propositions 11.2(iii) and (iv) in the final subsection, we will give the promised proof of Proposition 11.13. This will be proved using similar techniques as for Proposition 11.2(i); however, since we work in a region where some $\beta_n \equiv 1$ (and all other $\beta_{n'} \equiv 0$), we do not need to worry about ω_∞^n and thus our job will be easier.

We need one more technical lemma about superposition operators before giving the proof.

Lemma 11.14. *Fix $x \in \widehat{W}_j$ non-zero and $I \subseteq \mathbb{R}$. There is a smooth map*

$$\widehat{W}_j(I \times Y, x) \rightarrow \widehat{W}_j\left(I \times Y, \frac{x}{\|x\|_{L^2(Y)}}\right), \quad \gamma \mapsto \frac{\gamma}{\|\gamma\|_{L^2(Y)}},$$

defined in a neighborhood of any $\gamma \in \widehat{W}_j(I \times Y, x)$ with $\|\gamma(t)\|_{L^2(Y)} \geq \epsilon > 0$ for all $t \in I$. Further, this family of maps is smooth in x .

Proof. Without loss of generality, we assume that $\|x\|_{L^2(Y)} = 1$. (We can reduce to this case by dividing both γ and x by $\|x\|_{L^2(Y)}$.)

Let

$$\zeta := \gamma - x \in \widehat{W}_j(I \times Y).$$

We are interested in showing that the following quantity is in $\widehat{W}_j(I \times Y)$:

$$\begin{aligned} \frac{\gamma}{\|\gamma\|_{L^2(Y)}} - x &= \frac{\zeta + x}{\|\zeta + x\|_{L^2(Y)}} - x \\ &= \zeta \cdot (\|\zeta + x\|_{L^2(Y)}^{-1} - 1) + x \cdot (\|\zeta + x\|_{L^2(Y)}^{-1} - 1) + \zeta. \end{aligned}$$

Since $\zeta \in \widehat{W}_j(I \times Y)$ and $x \in \widehat{W}_j$, in view of Lemma 11.8, it suffices to show that

$$(259) \quad \|\zeta + x\|_{L^2(Y)}^{-1} - 1 \in L_j^2(I; \mathbb{R}).$$

Set

$$g := \|\zeta + x\|_{L^2(Y)}^2 - 1 = \langle \zeta, \zeta \rangle_{L^2(Y)} + 2\langle \zeta, x \rangle_{L^2(Y)}.$$

By applying Lemma 11.7, we see that $g \in L_j^2(I; \mathbb{R})$. The expression in (259) can be written as $h \circ g$, where $h(y) = (y+1)^{-1/2} - 1$. Note that $h(0) = 0$. Thus, we can apply Lemma 11.9(b), and deduce that $h \circ g \in L_j^2(I; \mathbb{R})$, as desired.

Smoothness with respect to γ follows from the smoothness statements in Lemmas 11.7, 11.8 and 11.9(b).

Similar arguments can be used to prove smoothness with respect to x . \square

With the above lemma, we can prove Proposition 11.13.

Proof of Proposition 11.13. Recall that by our assumptions on the path γ , we have that $\beta_i(\gamma) \equiv 1$ for some index i , and $\beta_n(\gamma) \equiv 0$ for all $n \neq i$. Writing $\gamma(t) = (a(t), s(t), \phi(t))$, we have

$$\Xi_\lambda(\gamma) = \left(a(t) + a_\infty^i - a_\lambda^i, s(t) + s_\infty^i - s_\lambda^i, \frac{\phi(t) + \langle \phi(t), \phi_\infty^i \rangle_{L^2(Y)} \cdot v_\lambda^i}{\|\phi(t) + \langle \phi(t), \phi_\infty^i \rangle_{L^2(Y)} \cdot v_\lambda^i\|_{L^2(Y)}} \right).$$

Recall that the spinorial component of $\Xi_\lambda(\gamma)$ in this case is written as $\frac{f_\lambda(\gamma(t))}{\|f_\lambda(\gamma(t))\|_{L^2(Y)}}$. The result now follows from applying Lemmas 11.7, 11.8 and 11.14 using the lower bounds given in (248). \square

11.4. Extensions of Ξ_λ to other path spaces. With the main technical results established, we are easily able to extend Ξ_λ to larger path spaces to complete the proof of Proposition 11.2 by proving parts (iii) and (iv).

Proposition 11.2(iii). Recall from Section 5.7 the space $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x, y)$. This space is larger than $\widetilde{W}_j^\tau(x, y)$ due to the condition of pseudo-temporal gauge as opposed to temporal gauge. Since $\widetilde{\mathcal{B}}_j^{\text{gC}, \tau}([x_\lambda], [y_\lambda])$ is the quotient of $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda)$ by the gauge group $\mathcal{G}_{j+1}^{\text{gC}}(\mathbb{R} \times Y)$, we will extend Ξ_λ to a diffeomorphism

$$\Xi_\lambda : \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda) \rightarrow \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\infty, y_\infty),$$

which commutes with the action by $\mathcal{G}_{j+1}^{\text{gC}}(\mathbb{R} \times Y)$, and this will induce the desired diffeomorphism for $\widetilde{\mathcal{B}}_j^{\text{gC}, \tau}$.

Let $(a(t) + \alpha(t)dt, s(t), \phi(t)) \in \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda)$, where $(a(t), s(t), \phi(t)) \in \widetilde{W}_j^\tau(x_\lambda, y_\lambda)$. By Proposition 11.2(ii), we have that $\Xi_\lambda(a(t), s(t), \phi(t)) \in \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\infty, y_\infty)$. We now define

$$\Xi_\lambda(a(t) + \alpha(t)dt, s(t), \phi(t)) = \Xi_\lambda(a(t), s(t), \phi(t)) + (\alpha(t)dt, 0, 0).$$

Since Ξ_λ induces a diffeomorphism from $\widetilde{W}_j^\tau(x_\lambda, y_\lambda)$ to $\widetilde{W}_j^\tau(x_\infty, y_\infty)$, it is clear that Ξ_λ induces a diffeomorphism from $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda)$ to $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\infty, y_\infty)$. Thus, it remains to see that this induced map respects the (four-dimensional) gauge action. Let $u \in \mathcal{G}_{j+1}^{\text{gC}, \tau}(\mathbb{R} \times Y)$. For notation, we will write $u(t) \cdot V(t)$ to mean the path obtained by applying $u(t)$ pointwise.

Using the S^1 -equivariance of $\Xi_\lambda : \widetilde{W}_j^\sigma \rightarrow \widetilde{W}_j^\sigma$, we have

$$\begin{aligned} u \cdot \Xi_\lambda(a(t) + \alpha(t)dt, s(t), \phi(t)) &= u \cdot (\Xi_\lambda(a(t), s(t), \phi(t)) + (\alpha(t)dt, 0, 0)) \\ &= (-u^{-1} \frac{du}{dt} dt, 0, 0) + u(t) \cdot \Xi_\lambda(a(t), s(t), \phi(t)) + (\alpha(t)dt, 0, 0) \\ &= (-u^{-1} \frac{du}{dt} dt, 0, 0) + \Xi_\lambda(a(t), s(t), u(t) \cdot \phi(t)) + (\alpha(t)dt, 0, 0) \\ &= \Xi_\lambda((-u^{-1} \frac{du}{dt} dt + \alpha(t)dt + a(t), s(t), u(t) \cdot \phi(t))) \\ &= \Xi_\lambda(u \cdot (a(t) + \alpha(t)dt, s(t), \phi(t))). \end{aligned}$$

It follows that Ξ_λ induces a diffeomorphism from $\widetilde{\mathcal{B}}_j^{\text{gC}, \tau}([x_\lambda], [y_\lambda])$ to $\widetilde{\mathcal{B}}_j^{\text{gC}, \tau}([x_\infty], [y_\infty])$. \square

Proposition 11.2(iv). We would like to show that Ξ_λ induces a bundle map

$$\begin{array}{ccc} \mathcal{V}_j^{\text{gC}, \tau} & \longrightarrow & \mathcal{V}_j^{\text{gC}, \tau} \\ \downarrow & & \downarrow \\ \widetilde{\mathcal{B}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda) & \xrightarrow{\Xi_\lambda} & \widetilde{\mathcal{B}}_j^{\text{gC}, \tau}(x_\infty, y_\infty). \end{array}$$

which is a diffeomorphism on the fibers.

Here, we recall that the bundle $\mathcal{V}_j^{\text{gC}, \tau}(\mathbb{R} \times Y)$ over $\widetilde{\mathcal{B}}_j^{\text{gC}, \tau}([x], [y])$ comes from quotienting the bundle $\mathcal{V}_j^{\text{gC}, \tau}(\mathbb{R} \times Y)$ over $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}([x], [y])$ by the gauge action. Therefore, it suffices to show that Ξ_λ induces a $\mathcal{G}_{j+1}^{\text{gC}, \tau}(\mathbb{R} \times Y)$ -equivariant bundle map

$$\begin{array}{ccc} \mathcal{V}_j^{\text{gC}, \tau} & \xrightarrow{\Xi_\lambda} & \mathcal{V}_j^{\text{gC}, \tau} \\ \downarrow & & \downarrow \\ \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\lambda, y_\lambda) & \xrightarrow{\Xi_\lambda} & \widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x_\infty, y_\infty), \end{array}$$

which is a diffeomorphism on the fibers.

Recall that the bundle $\mathcal{V}_j^{\text{gC}, \tau}$ over $\widetilde{\mathcal{C}}_j^{\text{gC}, \tau}(x, y)$ has fiber over $(a(t) + \alpha(t)dt, s(t), \phi(t))$ consisting of paths in $(b(t), r(t), \psi(t))$ satisfying $\text{Re}\langle \phi(t), \psi(t) \rangle_{L^2(Y)} = 0$ for all t . In other words, if we write $\gamma(t) = (a(t), s(t), \phi(t))$, then the fiber over $(a(t) + \alpha(t)dt, s(t), \phi(t))$ consists of paths η such that $\eta(t) \in \mathcal{T}_{j, \gamma(t)}^{\text{gC}, \sigma}$ for all t .

Since $\Xi_\lambda : \widetilde{W}_j^\sigma \rightarrow \widetilde{W}_j^\sigma$ is a diffeomorphism, we have that for any $x \in \widetilde{W}_j^\sigma$, we have an induced linear isomorphism

$$(260) \quad \mathcal{D}_x \Xi_\lambda : \mathcal{T}_{j,x}^{\text{gC},\sigma} \xrightarrow{\cong} \mathcal{T}_{j,\Xi_\lambda(x)}^{\text{gC},\sigma}.$$

Therefore, for an element $\eta \in \mathcal{V}_j^{\text{gC},\tau}$ which sits over $(a(t) + \alpha(t)dt, s(t), \phi(t))$, we can define $\Xi_\lambda(\eta)$ by pushforward. More precisely, $\Xi_\lambda(\eta)$ is the path in $\mathcal{V}_j^{\text{gC},\tau}$ given by

$$(\mathcal{D}_{(a(t),s(t),\phi(t))} \Xi_\lambda)(\eta(t)).$$

By construction $\Xi_\lambda(\eta)$ sits over $\Xi_\lambda(a(t) + \alpha(t)dt, s(t), \phi(t))$, and thus we have a bundle map. It is not difficult to check that $(\Xi_\lambda)_*$ induces a diffeomorphism of the fiber over $(a(t) + \alpha(t)dt, s(t), \phi(t))$, as $(\Xi_\lambda^{-1})_*$ provides the inverse. The gauge-equivariance follows from the gauge-equivariance of Ξ_λ on $\mathcal{C}^{\text{gC},\tau}$.

The analogous result for $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$ now follows from the above arguments together with the fact that if $[x_\infty] \neq [y_\infty]$, then the \mathbb{R} -action is free on $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\infty], [y_\infty])$ and $\widetilde{\mathcal{B}}_j^{\text{gC},\tau}([x_\lambda], [y_\lambda])$. \square

12. CONVERGENCE OF APPROXIMATE TRAJECTORIES

In this section and the subsequent one, we establish the analogous results of Section 7 for flow trajectories instead of stationary points. In this section, we focus on results related to the convergence of approximate trajectories to honest trajectories as $\lambda \rightarrow \infty$. We now return to the case that $\lambda = \lambda_i^\bullet \gg 0$; we will often omit this assumption from the discussion.

12.1. Convergence downstairs. We start by discussing trajectories of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ in the blow-down. There are two kinds of convergence results that one expects. The simpler one is C_{loc}^∞ convergence of parameterized trajectories, and the more refined one is convergence of unparameterized trajectories to a broken trajectory. We already know that the former kind holds:

Proposition 12.1. *Let $I \subseteq \mathbb{R}$ be a closed interval, and $\gamma_n : I \rightarrow W$ be a sequence of trajectories of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ contained in $B(2R)$, where $\lambda_n \rightarrow \infty$. Then, there exists a subsequence of γ_n for which the restrictions to any subinterval $I' \Subset I$ converge in the C^∞ topology of $W(I' \times Y)$ to γ , a trajectory of $\mathcal{X}_q^{\text{gC}}$.*

Proof. As discussed in Section 6.1, $l + c_q$ satisfies the hypotheses of Proposition 6.2. Thus, the result follows from Lemma 6.4. \square

In particular, when $I = \mathbb{R}$, the conclusion of Proposition 12.1 is convergence in the C_{loc}^∞ topology of $W(\mathbb{R} \times Y)$. We denote the resulting topological space by $W_{loc}(\mathbb{R} \times Y)$.

We now seek to show that a sequence of unparameterized trajectories for the approximate equations converge, in a certain sense, to a broken trajectory of $\mathcal{X}_q^{\text{gC}}$. It will be convenient to work in the quotient W/S^1 . Given $x \in W$, we write $[x]$ for its projection to W/S^1 . If x is a stationary point (of $\mathcal{X}_q^{\text{gC}}$ or $\mathcal{X}_{q^\lambda}^{\text{gC}}$), we will say that $[x]$ is a *stationary point class*.

Furthermore, given an interval $I \subset \mathbb{R}$ and a trajectory $\gamma : I \rightarrow W$ of $\mathcal{X}_q^{\text{gC}}$ or $\mathcal{X}_{q^\lambda}^{\text{gC}}$, we consider the associated *parameterized trajectory class*

$$[\gamma] \in W(I \times Y)/S^1,$$

where S^1 acts by constant gauge transformations. If $I = \mathbb{R}$, we can further divide by reparameterizations (translations in the domain \mathbb{R}), and obtain the *unparameterized trajectory class*

$$[\check{\gamma}] \in W(\mathbb{R} \times Y)/(S^1 \times \mathbb{R}).$$

When $I = \mathbb{R}$, Lemma 16.2.4 in [KM07], translated into slicewise Coulomb gauge, shows that any parameterized trajectory class $[\gamma]$ of $\mathcal{X}_q^{\text{gC}}$ with finite energy admits limit points $[x]$ and $[y]$ at $\pm\infty$, with $[x]$ and $[y]$ being stationary point classes of $\mathcal{X}_q^{\text{gC}}$. Corollary 8.19 gives the analogous result for trajectory classes of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, provided these come from trajectories contained in $B(2R)$. Since the limit points are unchanged by reparameterizations, we can also talk about the limit points of unparameterized trajectory classes.

Definition 12.2. *Fix $[x]$ and $[y]$, stationary point classes of $\mathcal{X}_q^{\text{gC}}$. An (unparameterized) broken trajectory class of $\mathcal{X}_q^{\text{gC}}$ from $[x]$ to $[y]$ consists of*

- an integer $m \geq 0$, the number of components;
- an $(m+1)$ -tuple of stationary point classes of $\mathcal{X}_q^{\text{gC}}$: $[x] = [x_0], \dots, [x_m] = [y]$
- an unparameterized trajectory class $[\check{\gamma}_i]$ of $\mathcal{X}_q^{\text{gC}}$ from $[x_{i-1}]$ to $[x_i]$, for every $i = 1, \dots, m$.

We will represent broken trajectory classes by the tuple $[\check{\gamma}] = ([\check{\gamma}_1], \dots, [\check{\gamma}_m])$.

Next, we want to say what it means for unparameterized trajectory classes of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ to converge to a broken trajectory class of $\mathcal{X}_q^{\text{gC}}$, as $\lambda_n \rightarrow \infty$. This is done in Definition 12.3 below, which is inspired by the construction of the topology on the space of broken trajectories of \mathcal{X}_q , in [KM07, Section 16.1]. Given a trajectory $\gamma : \mathbb{R} \rightarrow W$ and $s \in \mathbb{R}$, recall that we write $\tau_s \gamma$ for the translate, $(\tau_s \gamma)(t) = \gamma(s + t)$.

From Proposition 7.7 and Lemma 7.9, for any stationary point class $[x]$ of $\mathcal{X}_q^{\text{gC}}$, there is a corresponding (nearby) stationary point class of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, which we denote by $[x_\lambda]$, and similarly in the blow-up. We also recall from the discussion between Corollaries 8.20 and 8.21 that for $\lambda \gg 0$, trajectories $[\gamma_\lambda]$ of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$ in $(W^\lambda \cap B(2R)^\sigma)/S^1$ connecting stationary points $[x_\lambda]$ and $[y_\lambda]$ are contained in $\mathcal{B}_k^{\text{gC}, \tau}([x_\lambda], [y_\lambda])$. In particular, such trajectories satisfy $\lim_{t \rightarrow -\infty} [\tau_t^* \gamma_\lambda] = [\gamma_{x_\lambda}]$ and $\lim_{t \rightarrow -\infty} [\tau_t^* \gamma_\lambda] = [\gamma_{y_\lambda}]$ in $\mathcal{B}_{k, \text{loc}}^{\text{gC}, \tau}(\mathbb{R} \times Y)$ analogous to Definition 4.16. We have the analogous result in the blow-down as well. We will use these facts implicitly throughout the next two sections.

Definition 12.3. Fix $[x]$ and $[y]$, stationary point classes of $\mathcal{X}_q^{\text{gC}}$. For $\lambda_n \rightarrow \infty$, consider a sequence of unparameterized trajectory classes $[\check{\gamma}_n]$ of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$, coming from trajectories

$$\gamma_n : \mathbb{R} \rightarrow W^{\lambda_n} \cap B(2R)$$

and such that the endpoints of $[\check{\gamma}_n]$ are the stationary point classes $[x_{\lambda_n}]$ (respectively $[y_{\lambda_n}]$) of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$ that correspond to $[x]$ (respectively $[y]$).

We say that $[\check{\gamma}_n]$ converges to the broken trajectory class $[\check{\gamma}_\infty] = ([\check{\gamma}_{\infty, 1}], \dots, [\check{\gamma}_{\infty, m}])$ if, for each $i = 1, \dots, m$, there exist sequences of real numbers $(s_{n,i})_{n \geq 0}$ with

$$s_{n,1} < s_{n,2} < \dots < s_{n,m}$$

and

$$s_{n,i} - s_{n,i-1} \rightarrow \infty \quad \text{as } n \rightarrow \infty,$$

such that the translates $[\tau_{s_{n,i}} \gamma_n]$ converge to some representative $[\gamma_{\infty, i}]$ of $[\check{\gamma}_{\infty, i}]$ in the quotient topology of $W_{\text{loc}}(\mathbb{R} \times Y)/S^1$.

Proposition 12.4. Fix $[x]$ and $[y]$, stationary point classes of $\mathcal{X}_q^{\text{gC}}$. Fix a sequence $\lambda_n \rightarrow \infty$ and a sequence of unparameterized trajectory classes $[\check{\gamma}_n]$ of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}}$, going from $[x_{\lambda_n}]$ to $[y_{\lambda_n}]$, and such that the representatives γ_n of $[\check{\gamma}_n]$ are contained in $W^{\lambda_n} \cap B(2R)$. Then, there exists a subsequence of $[\check{\gamma}_n]$ that converges to a broken trajectory class $[\check{\gamma}_\infty]$ of $\mathcal{X}_q^{\text{gC}}$, in the sense of Definition 12.3.

Proof. The proof is essentially the same as [KM07, Proposition 16.2.1]. We provide an outline.

Since the energy $\mathcal{E}(\gamma)$ of a trajectory γ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ is unchanged by constant gauge transformations, we can talk about the energy $\mathcal{E}([\gamma])$ of the respective parameterized trajectory class. Similarly, since the functions F_λ constructed in Section 8 are S^1 -invariant, we can talk about the drop in F_λ for a parameterized trajectory class. In view of Corollary 8.21, the energy and the drop in F_λ are commensurable.

Fix a compact interval I , constant $C > 0$, and a neighborhood $U_{[x]}$ of each stationary point class $[x]$ of $\mathcal{X}_q^{\text{gC}}$, where $[x]$ is thought of as a constant trajectory class in $W(I \times Y)/S^1$.

Using Proposition 12.1, for any other compact interval I' , we can find $\epsilon > 0$ independent of λ such that if $[\gamma]$ is a trajectory class of $\mathcal{X}_{\mathfrak{q}\lambda}^{\text{gC}}$ with energy at most C , and with energy at most ϵ when restricted to I' , then $[\gamma]|_{I \times Y}$ is contained in some $U_{[x]}$. (Compare with Lemma 16.2.2 in [KM07].)

Next, observe that, since the trajectory classes $[\gamma_n]$ go from $[x_{\lambda_n}]$ to $[y_{\lambda_n}]$, and

$$[x_{\lambda_n}] \rightarrow [x], \quad [y_{\lambda_n}] \rightarrow [y], \quad F_{\lambda_n} \rightarrow \mathcal{L}_{\mathfrak{q}} \text{ as } n \rightarrow \infty,$$

we have that the drop in F_{λ_n} along $[\gamma_n]$ is bounded. Hence, the energy of $[\gamma_n]$ is bounded, by a constant K independent of n . With $\epsilon > 0$ chosen as in the previous paragraph, we find that, for each n , there are at most $2K/\epsilon$ integers p such that

$$\mathcal{E}([\tau_p \gamma_n]|_{[-1,1]}) = \mathcal{E}([\gamma_n]|_{[p-1,p+1]}) > \epsilon.$$

Therefore, for all other p , we must have $[\tau_p \gamma_n] \in U_{[x]}$ for some stationary point class $[x]$ of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$.

Starting from here, for each $n \gg 0$, we decompose \mathbb{R} into finitely many intervals $I_i^n = [a_i^n, b_i^n]$ of fixed length, and intervals

$$J_0^n = (-\infty, a_1^n], \quad J_i^n = [b_i^n, a_{i+1}^n], \quad J_m^n = [b_m^n, \infty),$$

with the length of each J_i^n going to infinity as $n \rightarrow \infty$, and the number m of intervals being independent of n . The restriction of $[\gamma_n]$ to each J_i^n lies near a stationary point class of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$, and these point classes provide the breaking points of the limiting broken trajectory class. By applying Proposition 12.1 to the restrictions of $[\gamma_n]$ to I_i^n , we can arrange that they are convergent in C_{loc}^∞ . This gives the required convergence to a broken trajectory class. \square

12.2. Convergence of parameterized trajectories in the blow-up. We now move to the blow-up W^σ . We are interested in showing that the trajectories of $\mathcal{X}_{\mathfrak{q}\lambda}^{\text{gC},\sigma}$ are close to those of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$, given appropriate control on the λ -spinorial energy. The goal of this subsection is to establish the following convergence result for parameterized trajectories. Before doing so, a quick notational remark. Sometimes we will be interested in studying the image of a path in W^σ/S^1 in the blow-down. When doing so, we will use the notation γ^τ for the path upstairs and γ for the blow-down.

Proposition 12.5. *Fix $\omega > 0$ and a compact interval $I = [t_1, t_2] \subset \mathbb{R}$. Consider a smaller interval $I_\epsilon = [t_1 + \epsilon, t_2 - \epsilon]$ for $\epsilon > 0$. Suppose that*

$$\gamma_n^\tau : I \rightarrow (W^{\lambda_n} \cap B(2R))^\sigma$$

is a sequence of trajectories of $\mathcal{X}_{\mathfrak{q}\lambda_n}^{\text{gC},\sigma}$, where $\lambda_n \rightarrow \infty$. Furthermore, suppose that at the ends of I_ϵ we have

$$\Lambda_{\mathfrak{q}\lambda_n}(\gamma_n^\tau(t_1 + \epsilon)) \leq \omega, \quad \Lambda_{\mathfrak{q}\lambda_n}(\gamma_n^\tau(t_2 - \epsilon)) \geq -\omega,$$

for all n . Then, there exists a subsequence of γ_n^τ for which the restrictions to any $I' \Subset I_\epsilon$ converge in the C^∞ topology of $W^\tau(I_\epsilon \times Y)$ to γ^τ , a trajectory of $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$.

Remark 12.6. The analogous compactness result for trajectories of $\mathcal{X}_{\mathfrak{q}}^\sigma$ is Theorem 10.9.2 in [KM07]. However, there are a few differences between the statement of Proposition 12.5 and [KM07, Theorem 10.9.2]. Of course, the main one is that our result deals with trajectories of the vector fields $\mathcal{X}_{\mathfrak{q}\lambda_n}^{\text{gC},\sigma}$, which approximate the vector field $\mathcal{X}_{\mathfrak{q}}^{\text{gC},\sigma}$ in Coulomb gauge.

Another difference is that in [KM07, Theorem 10.9.2] one requires a bound on the energy of trajectories, that is, on the drop of the perturbed Chern-Simons-Dirac functional \mathcal{L}_q ; this is unnecessary in our setting since our trajectories are assumed to be in $B(2R)^\sigma$, which automatically gives bounds on \mathcal{L}_q . Also, in [KM07, Theorem 10.9.2], the assumption is that q is a k -tame perturbation, and the conclusion is convergence in L^2_{k+1} . In our setting, q is tame (for all k), and we will use bootstrapping and a diagonalization argument for subsequences to get convergence in C^∞ . Finally, in [KM07, Theorem 10.9.2], the conclusion is convergence after gauge transformations. In our setting trajectories start in temporal and slicewise global Coulomb gauge; the remaining gauge consists only of constant transformations in S^1 . Since S^1 is compact, there is no need to change trajectories by gauge before they converge. Compare the remark after Corollary 5.1.8 on [KM07, p.110].

Before proving Proposition 12.5, we need a couple of lemmas. The first is the following unique continuation result, analogous to [KM07, Proposition 10.8.1]:

Lemma 12.7. *Let $\gamma(t) = (a(t), \phi(t))$ be a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ in $B(2R)$ for some $\lambda \in (0, \infty]$. If $\phi(t) = 0$ for some t , then $\phi(t) \equiv 0$.*

Proof. Since $\gamma(t)$ is a trajectory, we have

$$-\frac{d}{dt}\phi(t) = D\phi(t) + (p^\lambda c_q)^1(\gamma(t)).$$

Furthermore, $\gamma(t)$ is continuous and thus $(p^\lambda c_q)^1(\gamma(t))$ is a continuous path in $L^2_k(Y; \mathbb{S})$ as well. By [KM07, Lemma 7.1.3], it suffices to show that $\|(p^\lambda c_q)^1(\gamma(t))\|_{L^2} \leq C\|\phi(t)\|_{L^2}$ for all t . This follows from Lemma 6.13. \square

The second result we need is the analogue of [KM07, Lemma 10.9.1]:

Lemma 12.8. *There is a constant $C > 0$ such that, for any $\lambda \gg 0$, any interval $[t_1, t_2] \subseteq \mathbb{R}$, trajectory $\gamma^\tau : [t_1, t_2] \rightarrow B(2R)^\sigma$ of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$, and any $t \in [t_1, t_2]$, we have*

$$\frac{d}{dt}\Lambda_{q^\lambda}(\gamma^\tau(t)) \leq C \cdot \|\mathcal{X}_{q^\lambda}^{\text{gC}}(\gamma(t))\|_{L^2_k(Y)},$$

where γ is the projection of γ^τ in the blow-down.

Proof. This is similar to the proof of Lemma 10.9.1 in [KM07]. Note that the bound in [KM07] involved a function $\zeta(\gamma(t))$; in our case we can take this to be a constant, because we assumed that $\gamma(t) \subset B(2R)$. \square

Proof of Proposition 12.5. The argument follows that of Theorem 8.1.1 and Theorem 10.9.2 in [KM07] nearly verbatim. We give the argument for completeness.

Write $\gamma_n^\tau(t) = (a_n(t), s_n(t), \phi_n(t))$. First, by Proposition 12.1 we can find a subsequence of the blown-down sequence $\gamma_n(t) = (a_n(t), s_n(t), \phi_n(t))$ for which the restriction to any $I' \Subset I$ converges in all $L^2_j(I' \times Y)$ norms. Choose I' such that $I_\epsilon \Subset I'$. In particular, we have bounds (and convergence) for all $L^2_j(I' \times Y)$ norms on $a(t)$. Therefore, we will now focus on the convergence of $s_n(t)$ and $\phi_n(t)$.

By Lemma 12.7, each trajectory is either slicewise reducible or irreducible. After passing to a further subsequence, we can assume that γ_n^τ are all of the same type. We first consider the case that $\gamma_n^\tau(t)$ are irreducible for all n and t . If the limit of γ_n is irreducible, we are done using the result in the blow-down. Therefore, we assume that the limit is reducible for all t . It will be useful to simultaneously think of $s_n(t)\phi_n(t)$ as a four-dimensional spinor

$z_n \Phi_n$, where $z_n = \|s_n \phi_n\|_{L^2(I_\epsilon \times Y)}$ and $\|\Phi_n\|_{L^2(I_\epsilon \times Y)} = 1$, so $z_n \Phi_n = s_n \phi_n$. (We have $z_n \neq 0$ since $s_n(t) \neq 0$ for every $t \in I'$.) The subtle change from I' to I_ϵ when computing L^2 norms will be used shortly.

Since γ_n is an approximate trajectory in the blow-down, so is $e^{i\theta} \gamma_n$, and therefore we see that

$$-D^+(e^{i\theta} z_n \Phi_n) = (p^{\lambda_n} c_q)^1(a_n, e^{i\theta} z_n \Phi_n),$$

for all θ . (Here D^+ denotes the Dirac operator in four-dimensions for the trivial connection induced by A_0 .) By differentiating with respect to θ and evaluating at $\theta = 0$, we obtain

$$-D^+(z_n \Phi_n) = \mathcal{D}_{\gamma_n}(p^{\lambda_n} c_q)^1(0, z_n \Phi_n),$$

and thus by complex linearity,

$$-D^+(\Phi_n) = \mathcal{D}_{\gamma_n}(p^{\lambda_n} c_q)^1(0, \Phi_n).$$

By Lemma 6.12, the $L^2(I' \times Y)$ bounds on Φ_n and $L_k^2(I' \times Y)$ bounds on γ_n give that the right-hand side is $L^2(I' \times Y)$ bounded. By ellipticity, we see that Φ_n is L_1^2 -bounded on any interior cylinder. We can do further bootstrapping to obtain $L_j^2(I_\epsilon \times Y)$ bounds on the four-dimensional spinor Φ_n for all j . In particular, we can arrange for a further subsequence for which Φ_n converges in all $L_j^2(I_\epsilon \times Y)$ norms to a spinor Φ with $\|\Phi\|_{L^2(I_\epsilon \times Y)} = 1$.

By Lemma 12.8 and the bounds on $\Lambda_{q^{\lambda_n}}$ at the endpoints of I_ϵ , we see that $\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))$ is uniformly bounded in n and $t \in I_\epsilon$. Since $-\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t)) = \frac{d}{dt} \log s_n(t)$, it follows that there exists a constant $K > 0$ independent of n and $t \in I_\epsilon$ such that

$$s_n(t) \geq K \|s_n\|_{L^2(I_\epsilon)} = K z_n,$$

and thus

$$\|\Phi_n(t)\|_{L^2(Y)} \geq K > 0$$

for all t . Since Φ_n converges to Φ in all $L_j^2(I_\epsilon \times Y)$ norms (and thus uniformly in all $L_j^2(Y)$ norms), we see that $\phi_n(t) = \frac{\Phi_n(t)}{\|\Phi_n(t)\|_{L^2(Y)}}$ is bounded in all $L_j^2(I \times Y)$ norms by Lemma 11.14. It follows that we have the desired convergence for the ϕ_n .

We now focus on the convergence for s_n . Of course, rather than obtaining bounds on $L_j^2(I_\epsilon \times Y)$ norms as above, we could have used any intermediate cylinder with $I_\epsilon \subset I'' \Subset I'$. We in fact opt for I'' , since we will need to bootstrap again to get $L_j^2(I_\epsilon)$ bounds on s_n . By (165) and (166), since both $\gamma_n(t)$ and $\phi_n(t)$ are bounded in all $L_j^2(I'' \times Y)$ norms, we see that $\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))$ is bounded in all $L_j^2(I'')$ norms. Since $\gamma_n^\tau(t)$ is a trajectory of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}, \sigma}$, (165) shows that

$$(261) \quad -\frac{d}{dt} s_n(t) = -\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t)) s_n(t).$$

Note that for all n and $t \in I'$, $0 < s_n(t) \leq 2R$, so we obtain uniform bounds on $\frac{d}{dt} s_n(t)$ by the bounds on $s_n(t)$ and $\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))$. We may continue to bootstrap using (261) together with Sobolev multiplication to obtain bounds on s_n in all $L_j^2(I_\epsilon)$ norms, and consequently convergence in C^∞ . This completes the proof in the case that the γ_n^τ are irreducible.

Now, we consider the case that $\gamma_n^\tau(t)$ is reducible. As in the irreducible case we have $L_1^2(I')$ bounds on $\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))$, by Lemma 12.8 and the bounds at the endpoints. In this case, choose $s_n^*(t)$ to solve

$$-\frac{d}{dt} s_n^*(t) = \Lambda_{q^{\lambda_n}}(a_n(t), 0, \phi_n(t)) s_n^*(t),$$

where we ask that $0 < s_n^*(t) < M$ for all $t \in I$ and n . Since this is a one-dimensional ODE defined on a compact interval, we can easily arrange for such a solution. It is straightforward to verify that

$$-\left(\frac{d}{dt} + \mathcal{X}_{q^{\lambda_n}}^{\text{gC}}\right)^1 (a_n(t), s_n^*(t)\phi_n(t)) = \mathcal{D}_{(a_n(t), 0)}(p^{\lambda_n} c_q)^1(0, \phi_n(t)).$$

We can now repeat the same arguments as above to obtain the desired convergence of $\phi_n(t)$. \square

Corollary 12.9. *Fix $\omega > 0$ and a closed (possibly non-compact) interval $I \subseteq \mathbb{R}$. Suppose that $\gamma_n^\tau : I \rightarrow (W^{\lambda_n} \cap B(2R))^\sigma$ is a sequence of trajectories of $\mathcal{X}_{q^{\lambda_n}}^{\text{gC}, \sigma}$, where $\lambda_n \rightarrow \infty$. Furthermore, suppose that $|\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))| \leq \omega$ for all n and $t \in I$. Then, there exists a subsequence of γ_n^τ for which the restrictions to any subinterval $I' \Subset I$ converge in the C^∞ topology of $W^\tau(I' \times Y)$ to γ^τ , a trajectory of $\mathcal{X}_q^{\text{gC}, \sigma}$.*

Proof. This follows from Proposition 12.5, by a diagonalization argument. \square

When $I = \mathbb{R}$, we let $W_{loc}^\tau(\mathbb{R} \times Y)$ be $W^\tau(\mathbb{R} \times Y)$ with the C_{loc}^∞ topology. Then, the conclusion of Corollary 12.9 is convergence in $W_{loc}^\tau(\mathbb{R} \times Y)$.

12.3. Near-constant approximate trajectories. This subsection contains several technical results, analogous to the ones in Sections 13.4 and 13.5 of [KM07]. We will use these to study the moduli spaces of broken trajectories of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, leading up to Propositions 12.17 and 12.22 below.

We start by establishing the analogue of Lemma 13.4.4 in [KM07], which gives bounds on the distance of a trajectory from a constant trajectory on $[t_1, t_2] \times Y$ in terms of \mathcal{L}_q . Write $I = [t_1, t_2]$. We follow the notation of Sections 8.2 and 8.5. More precisely, given a stationary point $x_\infty = (a_\infty, \phi_\infty)$ of $\mathcal{X}_q^{\text{gC}}$, there is an associated orbit \mathcal{O}_λ of stationary points of $\mathcal{X}_{q^\lambda}^{\text{gC}}$. We define $x_\lambda = (a_\lambda, \phi_\lambda)$ to be the point in \mathcal{O}_λ which is L^2 closest to x_∞ .

Lemma 12.10. *Fix a stationary point $x_\infty = (a_\infty, \phi_\infty)$ of $\mathcal{X}_q^{\text{gC}}$, which we treat as a (constant) trajectory of $\mathcal{X}_q^{\text{gC}}$ on $I \times Y$. There exists an S^1 -invariant neighborhood U of x_∞ in $W_1(I \times Y)$ and a constant C independent of $\lambda \gg 0$ satisfying the following. If $\gamma \in U$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, there exists a constant gauge transformation $u \in S^1$ such that*

$$\|u \cdot \gamma - x_\lambda\|_{L_1^2(I \times Y)} \leq C(F_\lambda(\gamma(t_2)) - F_\lambda(\gamma(t_1))).$$

Proof. In [KM07, Lemma 13.4.4], Kronheimer and Mrowka choose a gauge transformation of γ that moves it into the Coulomb-Neumann slice through x_∞ . In our context, since we prefer to work with γ in slicewise global Coulomb gauge, we choose $u \in S^1$ so that

$$\text{Re}\langle u \cdot \gamma(t_1), (0, i\phi_\lambda) \rangle_{L^2(Y)} = 0.$$

To suppress u from the notation, for the rest of the proof, we assume $\text{Re}\langle \gamma(t_1), (0, i\phi_\lambda) \rangle_{L^2(Y)} = 0$, so that $u = 1$.

We have

$$(262) \quad \|\gamma - x_\lambda\|_{L_1^2(I \times Y)}^2 = \int_{t_1}^{t_2} \|\gamma - x_\lambda\|_{L_1^2(Y)}^2 dt + \int_{t_1}^{t_2} \left\| \frac{d}{dt} \gamma \right\|_{L^2(Y)}^2 dt.$$

In view of (214), it suffices to bound the right-hand side by a constant times $\int_{t_1}^{t_2} \|\frac{d}{dt}\gamma\|_{L^2(Y)}^2 dt$, or equivalently, $\int_{t_1}^{t_2} \|\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}(\gamma(t))\|_{L^2(Y)}^2 dt$. Since the second term in (262) is already of this form, it suffices to bound the first term, i.e., to find a constant C such that

$$\int_{t_1}^{t_2} \|\gamma - x_\lambda\|_{L_1^2(Y)}^2 dt \leq C \cdot \int_{t_1}^{t_2} \|\frac{d}{dt}\gamma\|_{L^2(Y)}^2 dt = C \cdot \int_{t_1}^{t_2} \|\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}(\gamma(t))\|_{L^2(Y)}^2 dt.$$

By our choice of \mathbf{q} , we have that x_∞ and x_λ are non-degenerate stationary points of $\mathcal{X}_{\mathbf{q}}^{\text{gC}}$ and $\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}$, respectively. There exist constants $C' > 0$ and $\delta > 0$ (independent of λ) such that for $\|x - x_\lambda\|_{L_1^2(Y)} < \delta$, we have

$$(263) \quad \|x - x_\lambda\|_{L_1^2(Y)}^2 \leq C' \left(\|\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}(x)\|_{L^2(Y)}^2 + |\text{Re}\langle x, (0, i\phi_\lambda) \rangle_{L^2(Y)}|^2 \right),$$

by an argument similar to that in Lemma 8.12, where we do not include the second term on the right-hand side if x_λ is reducible. Indeed, if x is real L^2 orthogonal to $(0, i\phi_\lambda)$, then the inequality follows directly from the non-degeneracy of $\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}$. More generally, we use

$$\begin{aligned} \|x - x_\lambda\|_{L_1^2(Y)} &\leq \|e^{i\theta}x - x_\lambda\|_{L_1^2(Y)} + \|x - e^{i\theta}x\|_{L_1^2(Y)} \\ &\leq C' \|\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}(x)\|_{L^2(Y)} + |(1 - e^{i\theta})| \|x\|_{L_1^2(Y)} \\ &\leq C' \|\mathcal{X}_{\mathbf{q}^\lambda}^{\text{gC}}(x)\|_{L^2(Y)} + C' |\text{Re}\langle x, (0, i\phi_\lambda) \rangle_{L^2(Y)}|, \end{aligned}$$

where θ is the angle between x and $(0, i\phi)$, i.e. $\sin \theta = \frac{\text{Re}\langle x, (0, i\phi_\lambda) \rangle_{L^2(Y)}}{\|x\|_{L^2(Y)} \|i\phi_\lambda\|_{L^2(Y)}}$. To see the last inequality, we use that x is L_1^2 bounded and that for irreducibles, the ϕ_λ is uniformly L^2 bounded above and below.

We therefore choose the neighborhood, U , in the statement of the lemma by extending a δ -neighborhood of x_λ in $L_1^2(I \times Y)$ to be gauge invariant.

By (263), for a trajectory $\gamma \in U$, we have the desired bounds in the case that x_λ is reducible. If x_λ is irreducible, it remains to bound

$$\int_{t_1}^{t_2} |\text{Re}\langle \gamma(t), (0, i\phi_\lambda) \rangle_{L^2(Y)}|^2 dt,$$

in terms of $\int_{t_1}^{t_2} \|\frac{d}{dt}\gamma\|_{L^2(Y)}^2 dt$. We write

$$\begin{aligned}
\int_{t_1}^{t_2} |\operatorname{Re}\langle \gamma(t), (0, i\phi_\lambda) \rangle_{L^2(Y)}|^2 dt &= \int_{t_1}^{t_2} |\operatorname{Re}\langle \gamma(t) - \gamma(t_1), (0, i\phi_\lambda) \rangle_{L^2(Y)}|^2 dt \\
&\leq C' \int_{t_1}^{t_2} \|\gamma(t) - \gamma(t_1)\|_{L^2(Y)}^2 dt \\
&\leq C' \int_{t_1}^{t_2} \left\| \int_{t_1}^t \left(\frac{d}{ds}\gamma\right) ds \right\|_{L^2(Y)}^2 dt \\
&\leq C' \int_{t_1}^{t_2} (t - t_1) \int_{t_1}^t \left\| \frac{d}{ds}\gamma \right\|_{L^2(Y)}^2 ds dt \\
&\leq C'(t_2 - t_1) \int_{t_1}^{t_2} \int_{t_1}^{t_2} \left\| \frac{d}{ds}\gamma \right\|_{L^2(Y)}^2 ds dt \\
&= C'(t_2 - t_1)^2 \int_{t_1}^{t_2} \left\| \frac{d}{ds}\gamma \right\|_{L^2(Y)}^2 ds,
\end{aligned}$$

where the fourth line follows from Cauchy-Schwarz. This completes the proof. \square

In [KM07, Sections 13.4 and 13.5], Lemma 13.4.4 is the starting point for a sequence of results about trajectories in neighborhoods of stationary points. In our setting, Lemma 12.10 above gives analogous results, by essentially the same arguments. We state the main results below, but omit the proofs.

First, by bootstrapping, we obtain the following from Lemma 12.10, analogous to [KM07, Proposition 13.4.7].

Lemma 12.11. *Fix a stationary point x_∞ of $\mathcal{X}_q^{\text{gC}}$, which we treat as a (constant) trajectory of $\mathcal{X}_q^{\text{gC}}$ on $I \times Y$. There exists an S^1 -invariant neighborhood U of x_∞ in $W_k(I \times Y)$ and a constant C independent of $\lambda \gg 0$ satisfying the following. If $\gamma \in U$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC}}$, there exists a constant gauge transformation $u \in S^1$ such that*

$$\|u \cdot \gamma - x_\lambda\|_{L_{k+1}^2(I' \times Y)} \leq C(F_\lambda(\gamma(t_2)) - F_\lambda(\gamma(t_1))),$$

for any compact interval $I' \Subset I$.

Next, we have two results about trajectories in the blow-up. These are the analogues of Proposition 13.4.1 and Corollary 13.4.8 in [KM07] respectively.

Proposition 12.12. *Let $x_\infty \in B(2R)^\sigma$ be a stationary point of $\mathcal{X}_q^{\text{gC}, \sigma}$, x_λ a nearby stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$, and $I' \Subset I = [t_1, t_2]$. Then, there exists a constant C and a gauge-invariant neighborhood U of the constant trajectory x_∞ in $W_k^\tau(I \times Y)$, independent of λ , such that for every trajectory $\gamma^\tau : I \rightarrow (W^\lambda \cap B(2R))^\sigma$ of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ which belongs to U , there is a gauge transformation $u \in S^1$ such that:*

(i) *if x_∞ is irreducible, then*

$$\|u \cdot \gamma^\tau - x_\lambda\|_{L_{k+1}^2(I' \times Y)} \leq C(F_\lambda(\gamma(t_1)) - F_\lambda(\gamma(t_2))),$$

(ii) *if x_∞ is reducible, then*

$$\|u \cdot \gamma^\tau - x_\lambda\|_{L_{k+1}^2(I' \times Y)} \leq C \left(\Lambda_{q^\lambda}(\gamma(t_1)) - \Lambda_{q^\lambda}(\gamma(t_2)) + (F_\lambda(\gamma(t_1)) - F_\lambda(\gamma(t_2)))^{\frac{1}{2}} \right).$$

Proposition 12.13. *Let $x_\infty \in B(2R)^\sigma$ be a stationary point of $\mathcal{X}_q^{\text{gC},\sigma}$ and x_λ the corresponding stationary point of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$. There is a constant C and a gauge-invariant neighborhood U of the constant trajectory in $W_k([-1, 1] \times Y)$ obtained from blowing down x_∞ , independent of $\lambda \gg 0$, with the following property. If $\gamma^\tau : [-1, 1] \rightarrow (W^\lambda \cap B(2R))^\sigma$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$ whose blow-down γ is in U , then*

$$\left. \frac{d}{dt} \Lambda_{q^\lambda}(\gamma^\tau(t)) \right|_{t=0} \leq C(F_\lambda(\gamma(-1)) - F_\lambda(\gamma(1)))^{\frac{1}{2}}.$$

For trajectories that converge to stationary points at the ends, we are able to obtain exponential decay on the value of F_λ near stationary points in the blow-up, analogous to that in [KM07, Proposition 13.5.1]. Since $F_\lambda : W^\lambda \cap B(2R) \rightarrow \mathbb{R}$ is S^1 -invariant, we obtain an induced map from $(W^\lambda \cap B(2R))^\sigma / S^1$ to \mathbb{R} . By a slight abuse of notation we will use F_λ for this induced map as well. Recall from Section 7, for every stationary point $[x_\infty]$ of $\mathcal{X}_q^{\text{agC},\sigma}$ with grading in $[-N, N]$, we have a unique corresponding stationary point $[x_\lambda]$ of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ for $\lambda \gg 0$.

Proposition 12.14. *Let $[x_\infty]$ be a stationary point of $\mathcal{X}_q^{\text{agC},\sigma}$ with grading in $[-N, N]$. There exists $\delta > 0$ such that for $\lambda \gg 0$, and for every trajectory $[\gamma] : [0, \infty) \rightarrow (W^\lambda \cap B(2R))^\sigma / S^1$ of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ with $\lim_{t \rightarrow \infty} [\tau_t^* \gamma] = [x_\lambda]$ in $L_{k,\text{loc}}^2$, there exists t_0 such that for $t \geq t_0$*

$$F_\lambda([\gamma(t)]) - F_\lambda([x_\lambda]) \leq C e^{-\delta t},$$

where $C = F_\lambda([\gamma(t_0)]) - F_\lambda([x_\lambda])$.

Here is a related result in the blow-down, which is the analogue of [KM07, Proposition 13.5.2].

Proposition 12.15. *Let x_∞ be a stationary point of $\mathcal{X}_q^{\text{gC}}$. There exists a neighborhood U of $[x_\infty]$ in $B(2R)/S^1$ and a constant $\delta > 0$ such that for $\lambda \gg 0$ and any trajectory $\gamma : [t_1, t_2] \rightarrow W^\lambda \cap U$ of $\mathcal{X}_{q^\lambda}^{\text{gC}}$ in $L_k^2([t_1, t_2] \times Y)$, we have the inequalities*

$$-C_2 e^{\delta(t-t_2)} \leq F_\lambda(\gamma(t)) - F_\lambda([x_\lambda]) \leq C_1 e^{-\delta(t-t_1)},$$

where

$$C_1 = |F_\lambda(\gamma(t_1)) - F_\lambda([x_\lambda])|, \quad C_2 = |F_\lambda(\gamma(t_2)) - F_\lambda([x_\lambda])|.$$

For a trajectory γ of $\mathcal{X}_{q^\lambda}^{\text{gC},\sigma}$, let us introduce the quantities

$$K_\lambda(\gamma) = \int_{\mathbb{R}} \left| \frac{d\Lambda_{q^\lambda}(\gamma)}{dt} \right| dt, \quad K_{\lambda,+}(\gamma) = \int_{\mathbb{R}} \left(\frac{d\Lambda_{q^\lambda}(\gamma)}{dt} \right)^+ dt,$$

which a priori may be infinite. Here, f^+ denotes $\max\{0, f\}$. Note that K_λ is finite if and only if $K_{\lambda,+}$ is finite. Further, if the result is finite, and γ is a trajectory from x_λ to y_λ , then

$$(264) \quad \Lambda_{q^\lambda}(x_\lambda) - \Lambda_{q^\lambda}(y_\lambda) = K_\lambda(\gamma) - 2K_{\lambda,+}(\gamma).$$

We will also use K_λ^I and $K_{\lambda,+}^I$ to restrict the domain of integration to any subinterval $I \subset \mathbb{R}$.

The exponential decay bound from Proposition 12.15, combined with Lemmas 12.8 and 12.11 give the following analogue of [KM07, Corollary 13.5.3].

Corollary 12.16. *Fix x_∞ a stationary point of $\mathcal{X}_q^{\text{gC}}$. Given a constant η , there is a gauge-invariant neighborhood $U \subset W_k([-1, 1] \times Y)$ of the constant trajectory x_∞ with the following property for $\lambda \gg 0$. Let $J \subset \mathbb{R}$ be any interval and $J' = J + [-1, 1]$. If we have a trajectory $\gamma^\tau : J' \rightarrow (W^\lambda \cap B(2R))^\sigma$ of $\mathcal{X}_{q^\lambda}^{\text{gC}, \sigma}$ such that $\tau_t \gamma$ are contained in U for all $t \in J$, then $K_{\lambda,+}^J(\gamma^\tau) \leq \eta$.*

12.4. Convergence of unparameterized trajectories in the blow-up. Our goal in this subsection is to prove analogues of the results in Section 12.1 in the blow-up. Similarly to the stationary point classes in the singular space W/S^1 that appeared Section 12.1, we will now consider stationary points $[x] \in W^\sigma/S^1$ of the vector fields $\mathcal{X}_q^{\text{agC}, \sigma}$ or $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$. We have *parameterized trajectories*

$$[\gamma] \in W^\tau(I \times Y)/S^1$$

of $\mathcal{X}_q^{\text{agC}, \sigma}$ or $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, and also, if $I = \mathbb{R}$, *unparameterized trajectories*

$$[\check{\gamma}] \in W^\tau(\mathbb{R} \times Y)/(S^1 \times \mathbb{R}).$$

Unlike in the blow-down, we omit the word “class” since the relevant objects arise from actual vector fields on the (smooth) quotients by S^1 .

We will only consider trajectories of $\mathcal{X}_q^{\text{agC}, \sigma}$ that limit to two stationary points. (Lemma 16.3.3 in [KM07], translated into slicewise Coulomb gauge, gives conditions for this to happen—but we will not need it here.) For trajectories of $\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, Corollary 8.20 shows the existence of limiting points, provided that the trajectories are contained in $(W^\lambda \cap B(2R))^\sigma/S^1$.

Finally, analogous to Definitions 12.2 and 12.3, we have the notions of unparameterized broken trajectories and convergence to them. In particular, convergence to a broken trajectory is defined in terms of the convergence of some parameterized representatives in $W_{loc}^\tau(\mathbb{R} \times Y)/S^1$.

The main focus of this subsection is the following.

Proposition 12.17. *Fix $[x]$ and $[y]$, stationary points of $\mathcal{X}_q^{\text{agC}, \sigma}$ with grading in $[-N, N]$. Fix $\lambda_n \rightarrow \infty$, and a sequence of unparameterized trajectories $[\check{\gamma}_n]$ of $\mathcal{X}_{q^{\lambda_n}}^{\text{agC}, \sigma}$, going from $[x_{\lambda_n}]$ to $[y_{\lambda_n}]$, and such that the representatives γ_n of $[\check{\gamma}_n]$ are contained in $(W^{\lambda_n} \cap B(2R))^\sigma$. Then, there exists a subsequence of $[\check{\gamma}_n]$ that converges to an unparameterized broken trajectory $[\check{\gamma}_\infty]$ of $\mathcal{X}_q^{\text{agC}, \sigma}$.*

Lemma 12.18. *Fix x and y stationary points of $\mathcal{X}_q^{\text{gC}, \sigma}$. There exists $C > 0$ such that for all $\lambda \gg 0$ and trajectories γ^τ from x_λ to y_λ , we have $K_\lambda(\gamma^\tau) \leq C$.*

Proof. This is analogous to [KM07, Lemma 16.3.1]. We outline the proof. We will show that any sequence γ_n^τ of trajectories from x_{λ_n} to y_{λ_n} has $K_{\lambda_n}(\gamma_n^\tau)$ bounded, as long as $\lambda_n \gg 0$. Note that since x_{λ_n} (respectively y_{λ_n}) are L_k^2 close to x (respectively y), it suffices to obtain uniform bounds on $K_{\lambda_n,+}(\gamma_n^\tau)$ by (264), since we have bounds on the λ_n -spinorial energies of x_{λ_n} and y_{λ_n} . Consider the projections γ_n of the trajectories γ_n^τ to the blow-down. Following the proof of Proposition 12.4, we find a decomposition of \mathbb{R} into a finite number, independent of n , of intervals of two types: J_n^i , on which γ_n is close to a constant trajectory as in Corollary 12.16, and I_n^i , which have a fixed length. Corollary 12.16 gives uniform bounds on $K_{\lambda_n,+}^{J_n^i}(\gamma_n^\tau)$. For the intervals I_n^i of fixed length, we can apply Lemma 12.8 to give

uniform upper bounds on $\frac{d}{dt}\Lambda_{q^{\lambda_n}}(\gamma_n^\tau(t))$. Since the lengths of the I_n^i are fixed, the values $K_{\lambda_n,+}^{I_n^i}$ are bounded. Thus, we obtain the desired bounds on $K_{\lambda_n}(\gamma_n^\tau)$. \square

Lemma 12.18 in fact gives uniform bounds on Λ_{q^λ} for approximate trajectories, which we now establish.

Lemma 12.19. *There exists a constant $C > 0$ independent of $\lambda \gg 0$ with the following property. If γ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ contained in $(W^\lambda \cap B(2R))^\sigma$, and going between two stationary points in the grading range $[-N, N]$, then $|\Lambda_{q^\lambda}(\gamma(t))| \leq C$ for all t .*

Proof. Since there are only finitely many stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ in the grading range $[-N, N]$, it suffices to establish the result for trajectories γ of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ between x_λ and y_λ for a fixed pair of stationary points. We have from (264) that

$$\begin{aligned} |\Lambda_{q^\lambda}(x_\lambda) - \Lambda_{q^\lambda}(\gamma(t))| &= |K_\lambda^{(-\infty,t]}(\gamma) - 2K_{\lambda,+}^{(-\infty,t]}(\gamma)| \\ &\leq 3K_\lambda(\gamma) \\ &\leq 3C, \end{aligned}$$

where C is the constant from Lemma 12.18. Since the x_λ are all L_k^2 close to the S^1 -orbit of a fixed stationary point x of $\mathcal{X}_q^{\text{agC}}$, we have that $|\Lambda_{q^\lambda}(x_\lambda)|$ is bounded independent of $\lambda \gg 0$. \square

The following is the analogue of [KM07, Lemma 16.3.2], and has a similar proof. It is a consequence of the convergence of parameterized trajectories in the blow-up, Corollary 12.9.

Lemma 12.20. *For each stationary point x of $\mathcal{X}_q^{\text{agC},\sigma}$, choose a gauge-invariant neighborhood $U_x \subset W_k^\tau(I \times Y)$ of the associated constant trajectory. Further, let I' be any other interval with non-zero length. Then, there exists $\epsilon > 0$ such that for $\lambda \gg 0$, if $\gamma \in M(y_\lambda, y'_\lambda)$ for stationary points y_λ, y'_λ of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ in the grading range $[-N, N]$ satisfies $K^{I'}(\gamma) \leq \epsilon$, then $\gamma|_{I \times Y} \in U_x$ for some x .*

Proof of Proposition 12.17. Lemma 12.19 guarantees that $\Lambda_{q^{\lambda_n}}$ is bounded for all trajectories between $[x_{\lambda_n}]$ and $[y_{\lambda_n}]$. Therefore, we may repeat the argument of Proposition 12.4, where we now apply Corollary 12.9 and Lemma 12.20 in place of Proposition 12.1 and we require that where we see bounds involving terms of the form $\mathcal{E}^I(\gamma)$, we also have analogous bounds on $K^I(\gamma)$. \square

Corollary 12.21. *Fix $\epsilon > 0$. For $\lambda \gg 0$, the following is true. Suppose that $[\gamma_\lambda]$ is a trajectory of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ from $[x_\lambda]$ to $[y_\lambda]$, such that either*

- (1) $[\gamma_\lambda]$ is not boundary-obstructed and $\text{gr}([x_\lambda], [y_\lambda]) = 1$ or
- (2) $[\gamma_\lambda]$ is boundary-obstructed and $\text{gr}([x_\lambda], [y_\lambda]) = 0$.

Furthermore, suppose that the gradings of $[x_\lambda]$ and $[y_\lambda]$ are in $[-N, N]$. Then, $[\gamma_\lambda]$ is ϵ -close in $W_{k,\text{loc}}(\mathbb{R} \times Y)/S^1$ to $[\gamma]$, a trajectory of $\mathcal{X}_q^{\text{agC},\sigma}$ with grading in $[-N, N]$.

Proof. We argue by contradiction. Suppose there exists a sequence of trajectories $[\gamma_{\lambda_n}]$, where $\lambda_n \rightarrow \infty$, which are always at least distance ϵ from trajectories of $\mathcal{X}_q^{\text{agC},\sigma}$. Without loss of generality, none of the trajectories are boundary-obstructed. A similar argument for the index 0 and boundary-obstructed case follows similarly. By Proposition 12.17, there exists a subsequence such that the unparameterized trajectories converge to a broken

trajectory $[\check{\gamma}_\infty]$ from $[x_\infty]$ to $[y_\infty]$, which is necessarily index 1 by Proposition 9.11. It remains to show that this trajectory is in fact unbroken.

We suppose that $[\check{\gamma}_\infty] = ([\check{\gamma}_{\infty,1}], \dots, [\check{\gamma}_{\infty,m}])$ with $m \geq 2$. By our non-degeneracy assumption, one $[\check{\gamma}_{\infty,i}]$ must be index 1, while the other trajectories must be index 0. Note that an index 0 trajectory must be boundary-obstructed (or else the moduli space is necessarily empty), and thus there can be at most one such trajectory. In particular, we see that $[\check{\gamma}_\infty] = ([\check{\gamma}_{\infty,1}], [\check{\gamma}_{\infty,2}])$. Since we are assuming that the trajectories $[\gamma_{\lambda_n}]$ are not boundary-obstructed, it follows that one of $[\check{\gamma}_{\infty,1}]$ or $[\check{\gamma}_{\infty,2}]$ goes from boundary-unstable to boundary-stable or one of $[x_\infty]$ or $[y_\infty]$ is irreducible. In either case, we have that that one of the components of $[\check{\gamma}_\infty]$ is an irreducible trajectory (see [KM07, Proposition 14.5.7] or the discussion at the end of Section 4.5). This implies that $M([x_\infty], [y_\infty])$ must contain an irreducible trajectory. This contradicts [KM07, Corollary 16.5.4], which implies that if $\check{M}([x_\infty], [y_\infty])$ contains a once-broken trajectory in the compactification and contains an irreducible trajectory, no component of the broken trajectory can be boundary-obstructed. \square

12.5. Convergence in L_k^2 . In Section 7, we used the inverse function theorem to find approximate stationary points near stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$. There we defined identifications Ξ_λ between $\mathfrak{C}_\mathcal{N}^\lambda$ and $\mathfrak{C}_\mathcal{N}$, which we extended to self-diffeomorphisms of W_j^σ in Lemma 11.1 and then to self-diffeomorphisms of path spaces in Proposition 11.2. We now use these self-diffeomorphisms to improve the $L_{k,\text{loc}}^2$ convergence from Proposition 12.17 to L_k^2 convergence.

Proposition 12.22. *Let $[\gamma_n] : \mathbb{R} \rightarrow (W^{\lambda_n} \cap B(2R))^\sigma / S^1$ be a sequence of trajectories of $\mathcal{X}_{q^{\lambda_n}}^{\text{agC},\sigma}$ between stationary points $[x_{\lambda_n}]$ and $[y_{\lambda_n}]$. Further, suppose that the unparameterized trajectory $[\check{\gamma}_n]$ converges to $[\check{\gamma}_\infty]$, an unbroken trajectory of $\mathcal{X}_q^{\text{agC},\sigma}$ from $[x_\infty]$ to $[y_\infty]$. Then, after possible reparameterization, $\Xi_{\lambda_n}([\gamma_n])$ converges to $[\gamma_\infty]$ in $\mathcal{B}_k^{\text{agC},\tau}([x_\infty], [y_\infty])$, where $[\gamma_\infty]$ is a representative of $[\check{\gamma}_\infty]$.*

Proof. We choose representative trajectories γ_n of $\mathcal{X}_{q^{\lambda_n}}^{\text{agC},\sigma}$ and γ_∞ of $\mathcal{X}_q^{\text{agC},\sigma}$ of the trajectory classes such that γ_n converges to γ_∞ . (This is easy to arrange, since these are only S^1 -equivalence classes.) This means there exists a sequence $s_n \in \mathbb{R}$ such that $\tau_{s_n} \gamma_n \rightarrow \gamma_\infty$ in $L_{k,\text{loc}}^2$. Since we may reparameterize $\Xi_{\lambda_n}([\gamma_n])$, we assume that no translations are necessary. Therefore, we are interested in showing that there exists a sequence of gauge transformations $u_n : \mathbb{R} \rightarrow S^1$ such that

$$\|u_n \cdot \Xi_{\lambda_n}(\gamma_n) - \gamma_\infty\|_{L_k^2(\mathbb{R} \times Y)} \rightarrow 0.$$

Fix $\epsilon > 0$. We will show that $\|\Xi_{\lambda_n}(\gamma_n) - \gamma_\infty\|_{L_k^2(\mathbb{R} \times Y)} < \epsilon$ for $n \gg 0$. Write x_∞ and y_∞ for the limit points of γ_∞ .

First, since $\gamma_n \rightarrow \gamma_\infty$ in $L_{k,\text{loc}}^2(\mathbb{R} \times Y)$, we also have that $\Xi_{\lambda_n}(\gamma_n) \rightarrow \gamma_\infty$ in $L_{k,\text{loc}}^2(\mathbb{R} \times Y)$ by Corollary 11.3. In particular, for any $T > 0$, we have for $n \gg 0$,

$$\|\Xi_{\lambda_n}(\gamma_n) - \gamma_\infty\|_{L_k^2([-T,T] \times Y)} < \epsilon/2.$$

Therefore, it suffices to show that we can find $T > 0$ such that for $n \gg 0$,

$$(265) \quad \|\Xi_{\lambda_n}(\gamma_n) - \gamma_\infty\|_{L_k^2([T,\infty) \times Y)} < \epsilon/4,$$

and likewise on $(-\infty, -T] \times Y$. We will establish the bounds for $[T, \infty) \times Y$; the case of $(-\infty, T] \times Y$ follows in the same way.

We first claim that there exists $T > 0$ such that for $n \gg 0$,

$$(266) \quad \|\Xi_{\lambda_n}(\gamma_n) - y_\infty\|_{L_k^2([T, \infty) \times Y)} < \epsilon/8.$$

By Corollary 11.4, we have that it suffices to establish

$$(267) \quad \|\gamma_n - y_{\lambda_n}\|_{L_k^2([T, \infty) \times Y)} \rightarrow 0, \quad n \rightarrow \infty.$$

This is an analogue of [KM07, Theorem 13.3.5], which implies that there exists a gauge transformation of γ_∞ which lands in $\mathcal{B}_k^{\text{gC}, \tau}([x_\infty], [y_\infty])$. However, we will explicitly keep track of the L_k^2 bounds and the necessary gauge transformation. First, by Proposition 12.12, there exists a constant gauge transformation $u_n^i \in S^1$ and a constant C , independent of i and λ , such that if y_∞ is reducible, then

$$(268) \quad \|u_n^i \cdot \gamma_n - y_{\lambda_n}\|_{L_k^2([i-2, i+2] \times Y)} \leq C \left((\Lambda_{q\lambda_n}(i-3) - \Lambda_{q\lambda_n}(i+3)) + (F_{\lambda_n}(i-3) - F_{\lambda_n}(i+3))^{\frac{1}{2}} \right),$$

and if y_∞ is irreducible, then

$$(269) \quad \|u_n^i \cdot \gamma_n - y_{\lambda_n}\|_{L_k^2([i-2, i+2] \times Y)} \leq C(F_{\lambda_n}(i-3) - F_{\lambda_n}(i+3)).$$

In Proposition 12.15 we have established the exponential decay of $F_{\lambda_n}(\gamma_n(t))$ independent of n . Further, by Proposition 12.13 and Proposition 12.15, we obtain uniform bounds on $\Lambda_{q\lambda_n}(i-3) - \Lambda_{q\lambda_n}(i+3)$ which decay exponentially as $i \rightarrow \infty$. We thus see that for any $T > 0$,

$$\sum_{i \geq T} \|u_n^i \cdot \gamma_n - y_{\lambda_n}\|_{L_k^2([i-2, i+2] \times Y)} \leq C(T)$$

where the constant $C(T) > 0$ depends only on T (i.e., independent of n) and converges to 0 as $T \rightarrow \infty$. To get the desired $L_k^2([T, \infty) \times Y)$ bounds on $\gamma_n - y_{\lambda_n}$, it suffices to prove that

$$\sum_{i \geq T} \|u_n^i \cdot y_{\lambda_n} - y_{\lambda_n}\|_{L_k^2([i-1/2, i+1/2] \times Y)} \leq C'(T),$$

for some constant $C'(T)$ that converges to 0 as $T \rightarrow \infty$. Without loss of generality, assume y_∞ is reducible (the irreducible case is similar). By (268), we have

$$\begin{aligned} & \| (u_n^i)^{-1} \cdot y_{\lambda_n} - (u_n^{i+1})^{-1} \cdot y_{\lambda_n} \|_{L_k^2([i-1/2, i+1/2] \times Y)} \\ & \leq \| (u_n^i)^{-1} \cdot y_{\lambda_n} - \gamma_n \|_{L_k^2([i-1/2, i+1/2] \times Y)} + \| (u_n^{i+1})^{-1} \cdot y_{\lambda_n} - \gamma_n \|_{L_k^2([i-1/2, i+1/2] \times Y)} \\ & = \| u_n^i \cdot \gamma_n - y_{\lambda_n} \|_{L_k^2([i-1/2, i+1/2] \times Y)} + \| u_n^{i+1} \cdot \gamma_n - y_{\lambda_n} \|_{L_k^2([i-1/2, i+1/2] \times Y)} \\ & \leq \| u_n^i \cdot \gamma_n - y_{\lambda_n} \|_{L_k^2([i-2, i+2] \times Y)} + \| u_n^{i+1} \cdot \gamma_n - y_{\lambda_n} \|_{L_k^2([i-1, i+3] \times Y)} \\ & \leq C \left((\Lambda_{q\lambda_n}(i-3) - \Lambda_{q\lambda_n}(i+3)) + ((\Lambda_{q\lambda_n}(i-2) - \Lambda_{q\lambda_n}(i+4)) \right. \\ & \quad \left. + (F_{\lambda_n}(i-3) - F_{\lambda_n}(i+3))^{\frac{1}{2}} + (F_{\lambda_n}(i-2) - F_{\lambda_n}(i+4))^{\frac{1}{2}} \right). \end{aligned}$$

The bounds from Proposition 12.13 and exponential decay of Proposition 12.15 show that

$$\sum_{i \geq T} \| (u_n^i)^{-1} \cdot y_{\lambda_n} - (u_n^{i+1})^{-1} \cdot y_{\lambda_n} \|_{L_k^2([i-1/2, i+1/2] \times Y)} \leq C'(T),$$

with $C'(T) \rightarrow 0$ as $T \rightarrow \infty$. We would like an analogous inequality without the inverses. Note that if u, u' are in S^1 , and $y = (a, s, \phi)$, we have that

$$u \cdot y - u' \cdot y = (0, 0, (u - u')\phi).$$

Since $|u - u'| = |u^{-1} - u'^{-1}|$, we see that

$$\|(u_n^i)^{-1} \cdot y_{\lambda_n} - (u_n^{i+1})^{-1} \cdot y_{\lambda_n}\|_{L_k^2([i-1/2, i+1/2] \times Y)} = \|u_n^i \cdot y_{\lambda_n} - u_n^{i+1} \cdot y_{\lambda_n}\|_{L_k^2([i-1/2, i+1/2] \times Y)}.$$

Therefore, we have

$$\sum_{i \geq T} \|u_n^i \cdot y_{\lambda_n} - u_n^{i+1} \cdot y_{\lambda_n}\|_{L_k^2([i-1/2, i+1/2] \times Y)} \leq C'(T).$$

In particular, this implies that as $i \rightarrow \infty$, u_n^i converges to some $u_n \in S^1$ (uniformly in n), and we see that

$$\sum_{i \geq T} \|u_n^i \cdot y_{\lambda_n} - u_n \cdot y_{\lambda_n}\|_{L_k^2([i-1/2, i+1/2] \times Y)} \leq C'(T).$$

Therefore, it remains to establish that $u_n = 1$ for each n . This follows since by the construction of u_n^i in Lemma 12.10, we have that $\text{Re}\langle u_n^i \cdot \gamma_n(i-2), (0, i\phi_n) \rangle_{L^2(Y)} = 0$ and $\lim_{t \rightarrow +\infty} \gamma_n(t) = y_{\lambda_n} = (a_n, \phi_n)$. We have now established (267) and thus (266) as claimed.

As mentioned above, [KM07, Theorem 13.3.5] gives $u : \mathbb{R} \rightarrow S^1$ such that $u \cdot \gamma_\infty - y_\infty$ is in $L_k^2([T, \infty) \times Y)$. In the above argument, we obtained an analogous result where we did not need the four-dimensional gauge transformation. By using the analogues of the results of this section for trajectories of $\mathcal{X}_q^{\text{gC}, \sigma}$ in place of $\mathcal{X}_{q_\lambda}^{\text{gC}, \sigma}$, we can repeat the above arguments to obtain the stronger statement for γ_∞ as well, i.e., that $\|\gamma_\infty - y_\infty\|_{L_k^2([T, \infty) \times Y)} < \infty$. (As explained at the beginning of the proof Lemma 12.10, we work in a different gauge slice instead of the Coulomb-Neumann slice, which allows the arguments of this section to pin down the gauge transformation precisely.) Therefore, for $T \gg 0$, we have that

$$\|\gamma_\infty - y_\infty\|_{L_k^2([T, \infty) \times Y)} < \epsilon/8.$$

Combining this inequality with (267) we obtain (265), which completes the proof. \square

We can use Proposition 12.17 to make more refined statements in the case of an approximate trajectory with small index.

Lemma 12.23. *Fix $\epsilon, N > 0$. For $\lambda \gg 0$, the following is true. Let $[\gamma_\lambda]$ be a trajectory from $[x_\lambda]$ to $[y_\lambda]$ such that either*

- (1) *$[\gamma_\lambda]$ is not boundary-obstructed and $\text{gr}([x_\lambda], [y_\lambda]) = 1$ or*
- (2) *$[\gamma_\lambda]$ is boundary-obstructed and $\text{gr}([x_\lambda], [y_\lambda]) = 0$.*

Furthermore, suppose that the gradings of $[x_\lambda]$ and $[y_\lambda]$ are in $[-N, N]$. Then $\Xi_\lambda([\gamma_\lambda])$ is ϵ -close in $L_k^2(\mathbb{R} \times Y)$ of $[\gamma]$, a trajectory of $\mathcal{X}_q^{\text{agC}, \sigma}$ from $[x_\infty]$ to $[y_\infty]$ with grading in $[-N, N]$.

Proof. The argument is the same as for Corollary 12.21, except after applying Proposition 12.17 to extract a convergent subsequence in $L_{k, \text{loc}}^2$, we improve this to L_k^2 convergence (after composition with Ξ_λ) using Proposition 12.22. \square

13. CHARACTERIZATION OF APPROXIMATE TRAJECTORIES

In this section, using the inverse function theorem, we will identify the moduli spaces of isolated approximate trajectories (in a fixed grading range) with those of actual trajectories. We will also produce similar identifications for the cut moduli spaces used to define the U -actions in Morse and Floer homology. Furthermore, we will show that all these identifications can be chosen to preserve orientations.

13.1. Stability. Let $[x_\infty]$ and $[y_\infty]$ be two stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$, with grading in $[-N, N]$; in other words, we have $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$. We assume that their grading difference is one and that we are not in the boundary-obstructed case. (Recall that boundary-obstructed means that $[x_\infty]$ is boundary-stable and $[y_\infty]$ is boundary-unstable.) Since q is an admissible perturbation, the moduli space

$$\check{M}^{\text{agC}}([x_\infty], [y_\infty]) = M^{\text{agC}}([x_\infty], [y_\infty])/\mathbb{R}$$

is a finite set of points. We can view $\check{M}^{\text{agC}}([x_\infty], [y_\infty])$ as the zero set of the section induced by $\mathcal{F}_q^{\text{gC},\tau}$ on the bundle $\mathcal{V}^{\text{gC},\tau}(\mathbb{R} \times Y)$ over $\check{\mathcal{B}}^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$, restricted to $\mathcal{B}^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$.

Consider the nearby stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$

$$[x_\lambda] = \Xi_\lambda^{-1}([x_\infty]), \quad [y_\lambda] = \Xi_\lambda^{-1}([y_\infty]).$$

We also have a moduli space of approximate Seiberg-Witten trajectories in $(B(2R) \cap W^\lambda)^\sigma$:

$$\check{M}^{\text{agC}}([x_\lambda], [y_\lambda]) = M^{\text{agC}}([x_\lambda], [y_\lambda])/\mathbb{R}.$$

By Corollary 9.3, the relative grading between $[x_\lambda]$ and $[y_\lambda]$ is also one. Further, by Proposition 8.1 and 10.2, for $\lambda = \lambda_i^\bullet$ with $i \gg 0$, $\check{M}^{\text{agC}}([x_\lambda], [y_\lambda])$ is regular and consists of finitely many points. Our main goal in this subsection is to prove:

Proposition 13.1. *For $\lambda = \lambda_i^\bullet \gg 0$, there is a one-to-one correspondence between the moduli spaces $\check{M}^{\text{agC}}([x_\infty], [y_\infty])$ and $\check{M}^{\text{agC}}([x_\lambda], [y_\lambda])$.*

The proof has two parts. First, we show that in a fixed neighborhood of any $[\gamma] \in \check{M}^{\text{agC}}([x_\infty], [y_\infty])$, there is a unique approximate trajectory. Second, we show that no other approximate trajectories between $[x_\lambda]$ and $[y_\lambda]$ exist.

We start with the first step. This is a stability result, similar to what we established in Proposition 7.7 for stationary points.

Proposition 13.2. *Suppose $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$ have $\text{gr}([x_\infty], [y_\infty]) = 1$, and we are not in the boundary-obstructed case. Fix a trajectory $[\gamma_\infty] \in M^{\text{agC}}([x_\infty], [y_\infty])$, and a small neighborhood U of $[\gamma_\infty]$ in $\mathcal{B}_k^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$.*

Then, for $\lambda \gg 0$, there exists a unique trajectory $[\gamma_\lambda] \in \check{M}^{\text{agC}}([x_\lambda], [y_\lambda])$ such that $\Xi_\lambda([\gamma_\lambda]) \in U$.

Proof. Recall that, because we do the blow-up construction in each time slice, the space $\mathcal{B}_k^{\text{gC},\tau}([x_\infty], [y_\infty])$ is not even a Banach manifold with boundary. However, it is a subspace of the Banach manifold $\check{\mathcal{B}}_k^{\text{gC},\tau}([x_\infty], [y_\infty])$, which consists of (equivalence classes of) paths $(a(t) + \alpha(t)dt, s(t), \phi(t))$, with $s(t)$ allowed to vary in all of \mathbb{R} .

By Proposition 11.2(iv), the map Ξ_λ gives rise to a bundle map $(\Xi_\lambda)_*$ that is part of a commutative diagram

$$\begin{array}{ccc} \mathcal{V}_k^{\text{gC},\tau}(\mathbb{R} \times Y) & \xrightarrow{(\Xi_\lambda)_*} & \mathcal{V}_k^{\text{gC},\tau}(\mathbb{R} \times Y) \\ \downarrow & & \downarrow \\ \tilde{\mathcal{B}}_k^{\text{gC},\tau}([x_\lambda], [y_\lambda])/\mathbb{R} & \xrightarrow{\Xi_\lambda} & \tilde{\mathcal{B}}_k^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}. \end{array}$$

We define a map

$$S : \tilde{\mathcal{B}}_k^{\text{gC},\tau}([x_\infty], [y_\infty]) \times (-1, 1) \rightarrow \mathcal{V}_k^{\text{gC},\tau}(\mathbb{R} \times Y) \times \mathbb{R}$$

by the formula

$$S([\gamma], r) = \left((\Xi_\lambda)_* \circ \mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau} \circ \Xi_\lambda^{-1}([\gamma]), r \right),$$

where $\mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau}$ is as in (167) and we wrote $\lambda = f^{-1}(|r|)$, with $f : (0, \infty] \rightarrow [0, 1)$ being the homeomorphism from Section 7.2.

By Proposition 11.2, we get that the section S is differentiable. Observe also that

$$S([\gamma], 0) = (\mathcal{F}_{\mathfrak{q}}^{\text{gC},\tau}([\gamma]), 0),$$

because Ξ_∞ is the identity. Therefore, the derivative of S at $([\gamma_\infty], 0)$ can be written in block form as

$$\begin{pmatrix} \mathcal{D}_{[\gamma_\infty]}^\tau \mathcal{F}_{\mathfrak{q}}^{\text{gC},\tau} & * \\ 0 & I \end{pmatrix}.$$

By our assumptions on \mathfrak{q} and Proposition 5.42, we get that $\mathcal{D}_{[\gamma_\infty]}^\tau \mathcal{F}_{\mathfrak{q}}^{\text{gC},\tau}$ is surjective. It is Fredholm index one when viewed as a map with domain $\mathcal{K}_{k,\gamma_\infty}^\tau$, the tangent space to $\tilde{\mathcal{B}}_k^{\text{gC},\tau}([x_\infty], [y_\infty])$ at $[\gamma_\infty]$, and range $\mathcal{K}_{k-1,\gamma_\infty}^{\text{gC},\tau}$. In our setting we divided by \mathbb{R} , so it becomes a map of index zero. Since it is surjective, it must be invertible. Given the block form above, it follows that $\mathcal{D}_{([\gamma_\infty], 0)}^\tau S$ is also invertible.

We now apply the inverse function theorem. For $r > 0$ small (that is, for $\lambda \gg 0$), we obtain a unique solution $[\gamma] \in U$ to the equation $S([\gamma], r) = (0, r)$. We then let

$$[\gamma_\lambda] = \Xi_\lambda^{-1}([\gamma]).$$

Because Ξ_λ is a diffeomorphism, we see that $\mathcal{F}_{\mathfrak{q}^\lambda}^{\text{gC},\tau}([\gamma_\lambda]) = 0$. Since $[\gamma_\lambda]$ has a temporal gauge representative in $\mathcal{C}_k^{\text{gC},\tau}(x_\lambda, y_\lambda)$ for some representatives x_λ, y_λ (see Remark 5.34), we see that $[\gamma_\lambda]$ gives the desired trajectory from $[x_\lambda]$ to $[y_\lambda]$.

Let γ_λ be a temporal gauge representative of $[\gamma_\lambda]$. We need to check that it actually takes values in $(B(2R) \cap W^\lambda)^\sigma$. We can repeat the argument above with the Sobolev coefficient $k+1$ instead of k . The resulting γ_λ in this case must be the same due to the uniqueness guaranteed by the implicit function theorem. Since we use the Sobolev coefficient $k+1$, we have that $\Xi_\lambda([\gamma_\lambda])$ must converge to $[\gamma]$ in an $L_{k+1}^2(\mathbb{R} \times Y)$ neighborhood. Since any temporal gauge representative of γ is contained in $B(R)^\sigma$, we see that $\Xi_\lambda(\gamma_\lambda)$ is contained in $B(3R/2)^\sigma$ for $\lambda \gg 0$. It is easy to see from the construction of Ξ_λ in Section 11 that γ_λ must be contained in $B(2R)^\sigma$. Further, we see that the values of $\Lambda_{\mathfrak{q}^\lambda}(\gamma_\lambda)$ are uniformly close to those of $\Lambda_{\mathfrak{q}}(\gamma)$, and thus uniformly bounded. By Lemma 6.15, we see that γ_λ is contained in $(W^\lambda)^\sigma$ for $\lambda \gg 0$, which completes the proof. \square

Remark 13.3. If the trajectory $[\gamma_\infty]$ from Proposition 13.2 is contained in the reducible locus, then so is the nearby approximate trajectory $[\gamma_\lambda]$. This follows from the invariance of the constructions in the proof under the obvious involutions on $\tilde{\mathcal{B}}^{\text{gC},\tau}([x_\lambda], [y_\lambda])$ and $\mathcal{V}_k^{\text{gC},\tau}(\mathbb{R} \times Y)$.

We are now ready to complete the proof of Proposition 13.1 with the second step, which is to show that there are no other approximate trajectories between $[x_\lambda]$ and $[y_\lambda]$.

Proof of Proposition 13.1. Let

$$\check{M}^{\text{agC}}([x_\infty], [y_\infty]) = \{[\gamma_\infty^1], \dots, [\gamma_\infty^m]\}.$$

For each $\ell = 1, \dots, m$, Proposition 13.2 guarantees the existence of a neighborhood U^ℓ of $[\gamma_\infty^\ell]$ and a unique approximate trajectory $[\gamma_\lambda^\ell]$ with $\Xi_\lambda([\gamma_\lambda^\ell]) \in U^\ell$.

We need to check that there are no other approximate trajectories in $\check{M}^{\text{agC}}([x_\lambda], [y_\lambda])$. Assume we had sequences $\lambda_n \rightarrow \infty$ and $[\gamma_n] \in \check{M}^{\text{agC}}([x_{\lambda_n}], [y_{\lambda_n}])$ such that $[\gamma_n]$ is not of the form $[\gamma_{\lambda_n}^\ell]$ for any n and ℓ . By assumption, each λ_n is some λ_i^\bullet . By applying Proposition 12.22 we can find a subsequence of the $[\gamma_n]$ such that $\Xi_{\lambda_n}([\gamma_n])$ converge to some $[\gamma_\infty^\ell]$ in L_k^2 . (We require that the λ_n be of the form λ_i^\bullet since we are invoking results of Section 12, which used this assumption throughout.) This means that the $\Xi_{\lambda_n}([\gamma_n])$ live inside U^ℓ for large n , and we obtain a contradiction with the uniqueness statement in Proposition 13.2. \square

We have an analogous result for the boundary-obstructed case.

Proposition 13.4. *Let $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$ be in the boundary-obstructed case, and suppose $\text{gr}([x_\infty], [y_\infty]) = 0$. For $\lambda = \lambda_i^\bullet \gg 0$, there is a one-to-one correspondence between $\check{M}^{\text{agC},\text{red}}([x_\infty], [y_\infty])$ and $\check{M}^{\text{agC},\text{red}}([x_\lambda], [y_\lambda])$.*

Proof. The proof is similar to that of Proposition 13.1, but with an application of the inverse function theorem in the subspace of $\tilde{B}_k^{\text{gC},\tau}([x_\infty], [y_\infty])/\mathbb{R}$ consisting of reducible paths. In this subspace, the corresponding linear operator will again be index zero, and in fact invertible. \square

13.2. The U -maps. In Section 4.7 we gave the definition of the U -maps in monopole Floer homology, in terms of cut-down moduli spaces $M([x], [y]) \cap \mathcal{Z}$ and $M^{\text{red}}([x], [y]) \cap \mathcal{Z}$; cf. Equations (55) and (56). Moreover, in Section 5.12 we identified these cut-down moduli spaces with similar ones in Coulomb gauge:

$$(270) \quad M^{\text{agC}}([x], [y]) \cap \mathcal{Z}^{\text{agC}} \text{ and } M^{\text{agC},\text{red}}([x], [y]) \cap \mathcal{Z}^{\text{agC}}.$$

Here, \mathcal{Z}^{agC} is the zero set of a section ζ^{agC} of the natural complex line bundle $E^{\text{agC},\sigma}$ over $W_k^\sigma/S^1 \subset W_{k-1/2}^\sigma/S^1$. The section is chosen so that the intersections in (270) are transverse.

We would like to further identify the moduli spaces from (270) with those consisting of approximate trajectories, that appear in the construction of equivariant Morse homology in finite dimensions; see Section 2.7. To obtain the cut-down moduli spaces in the approximations $(W^\lambda)^\sigma/S^1$, we simply consider the restriction of ζ^{agC} to those spaces.

Recall from Section 10 that we chose the perturbation \mathfrak{q} such that, for all $\lambda \in \{\lambda_1^\bullet, \lambda_2^\bullet, \dots\}$ sufficiently large, all the moduli spaces of flow lines for $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC},\sigma}$ on $(B(2R) \cap W^\lambda)^\sigma/S^1$ are regular. Since there is a countable number of such moduli spaces, we can choose the section ζ^{agC} such that it intersects all these moduli spaces (including the non-approximate ones)

transversely. This implies that we can define the U -maps on the corresponding Morse homology groups as in Section 2.7.

We now focus on the cut-down moduli spaces of the form

$$M^{\text{agC}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}} \text{ and } M^{\text{agC,red}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}},$$

for $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$. These moduli spaces suffice to determine the U -action in gradings from $-N$ to N .

Proposition 13.5. (a) *Suppose $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$ have $\text{gr}([x_\infty], [y_\infty]) = 2$ and are not boundary obstructed. For $\lambda = \lambda_i^\bullet \gg 0$, there is a one-to-one correspondence between $M^{\text{agC}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}$ and $M^{\text{agC}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}}$.*

(b) *Let $[x_\infty], [y_\infty] \in \mathfrak{C}_\mathcal{N}$ be reducible stationary points with $\text{gr}([x_\infty], [y_\infty]) = 1$ and boundary obstructed. For $\lambda = \lambda_i^\bullet \gg 0$, there is a one-to-one correspondence between $M^{\text{agC,red}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}$ and $M^{\text{agC,red}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}}$.*

Proof. Part (a) is the analogue of Proposition 13.1, and its proof is similar. Since the grading difference between $[x_\infty]$ and $[y_\infty]$ (and hence also between $[x_\lambda]$ and $[y_\lambda]$) is two, the cut-down moduli spaces in question are zero-dimensional. Consider a trajectory

$$[\gamma] \in M^{\text{agC}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}.$$

Recall that the intersection is taken at time $t = 0$, that is, we have $[\gamma(0)] \in \mathcal{Z}^{\text{agC}}$. The fact that the moduli space $M^{\text{agC}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}$ is cut out transversely at $[\gamma]$ can be rephrased in terms of the invertibility of the linear operator

$$\mathcal{D}_{[\gamma]}^\tau \mathcal{F}_{\mathfrak{q}}^{\text{gC}, \tau} \oplus \mathcal{D}_{[\gamma(0)]} \zeta^{\text{agC}} : \mathcal{T}_{k, [\gamma]}^{\text{gC}, \tau}([x_\infty], [y_\infty]) \rightarrow \mathcal{V}_{k-1, [\gamma]}^{\text{gC}, \tau}(Z) \oplus E_{[\gamma(0)]}^{\text{agC}, \sigma}.$$

An application of the inverse function theorem shows that in a neighborhood of $[\gamma]$ there is a unique element of $M^{\text{agC}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}}$ (after applying the corresponding diffeomorphism Ξ_λ). This produces at least as many elements of $M^{\text{agC}}([x_\lambda], [y_\lambda]) \cap \mathcal{Z}^{\text{agC}}$ as there are in $M^{\text{agC}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}$. To see that there are no more, notice that, given a sequence $[\gamma_n] \in M^{\text{agC}}([x_{\lambda_n}], [y_{\lambda_n}]) \cap \mathcal{Z}^{\text{agC}}$ with $\lambda_n \rightarrow \infty$, the sequence $\Xi_{\lambda_n}([\gamma_n])$ admits a convergent subsequence, and the limit must be an element of $M^{\text{agC}}([x_\infty], [y_\infty]) \cap \mathcal{Z}^{\text{agC}}$.

Part (b) is the analogue of Proposition 13.4, and again the proof is similar. \square

13.3. Orientations. Recall that in Section 5.13 we oriented the moduli spaces of trajectories of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}, \sigma}$, using an orientation data system σ^{gC} in Coulomb gauge. This was based on trivializing the determinant lines $\det(P_\gamma^{\text{gC}})$ for the operators P_γ^{gC} from (160). The same procedure, but using operators P_γ^{gC} defined with the perturbation \mathfrak{q}^λ instead of \mathfrak{q} , can be used to orient the moduli spaces of trajectories of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}, \sigma}$ (in infinite dimensions).

We would like to relate the orientations of the moduli spaces of trajectories of $\mathcal{X}_{\mathfrak{q}^\lambda}^{\text{gC}, \sigma}$ to the orientations in (finite dimensional) Morse homology and to the orientations of the moduli spaces of trajectories of $\mathcal{X}_{\mathfrak{q}}^{\text{gC}, \sigma}$ used in the definition of monopole Floer homology (recast in Coulomb gauge). For the moduli spaces of approximate Seiberg-Witten trajectories in the setting of finite dimensional Morse homology, we use the third construction of orientations discussed in Section 2.2, in terms of a specialized coherent orientation (cf. Definition 2.3). This is based on trivializing determinant lines $\det(\tilde{L}_\gamma)$, for the operators \tilde{L}_γ defined in (7).

To relate the operators P_γ^{gC} (with perturbation \mathfrak{q}^λ) and \tilde{L}_γ , we proceed as in the proof of Proposition 9.4. Indeed, P_γ^{gC} is the analogue of the operator Q_γ^{gC} considered in that

proposition, and \tilde{L}_γ is the analogue of (220); in both cases, what is different here is that we work on compact cylinders and add spectral projections on the boundaries. Recall from the proof of Proposition 9.4 that we have an orthogonal splitting $Q_\gamma^{\text{gC}} = \tilde{L}_\gamma \oplus F_\gamma$, where $F_\gamma : L_j^2(\mathbb{R}; (W^\lambda)^\perp) \rightarrow L_{j-1}^2(\mathbb{R}; (W^\lambda)^\perp)$ is defined by

$$F_\gamma(b(t), \psi(t)) = \left(\frac{d}{dt}b(t) + *d(b(t)), \frac{d}{dt}\psi(t) + D\psi(t) - \langle \phi(t), D\phi(t) \rangle_{L^2} \psi(t) \right).$$

Write \tilde{F}_γ for the analogous operator on the compact cylinder coupled with spectral projections. Note that while the operator \tilde{F}_γ depends on the path, the domain and target do not. We also write \tilde{F}_0 for the analogous operator without the term $\langle \phi(t), D\phi(t) \rangle_{L^2} \psi(t)$.

We seek to show that there are trivializations of $\det(\tilde{F}_\gamma)$ which respect concatenations of paths. First, note that the spectral decomposition of $(W^\lambda)^\perp$ corresponding to \tilde{F}_γ is independent of γ ; indeed, since $\phi(t) \in W^\lambda$ and $\psi(t) \in (W^\lambda)^\perp$, we have that $\|\langle \phi(t), D\phi(t) \rangle_{L^2} \psi(t)\|_{L^2(Y)} < \lambda \|\psi(t)\|_{L^2(Y)}$, while $\|D\psi(t)\|_{L^2(Y)} \geq \lambda \|\psi(t)\|_{L^2(Y)}$ and the claim follows. Further, we have the same spectral decomposition for any convex combination of \tilde{F}_γ and \tilde{F}_0 . Note that $(1-r)\tilde{F}_0 + r\tilde{F}_\gamma$ differs from \tilde{F}_0 by a compact operator. Thus, we have a homotopy between \tilde{F}_γ and \tilde{F}_0 through Fredholm operators, which induces an identification between $\det(\tilde{F}_\gamma)$ and $\det(\tilde{F}_0)$. This identification can easily be seen to respect concatenations. Therefore, we have an identification of $\det(\tilde{F}_\gamma)$, for each γ , with a fixed determinant line bundle (independent of γ) which respects concatenation. Thus, we can find compatible trivializations of $\det(\tilde{F}_\gamma)$ by simply fixing a trivialization of $\det(\tilde{F}_0)$.

The above discussion shows that a trivialization of $\det(P_\gamma^{\text{gC}})$ is equivalent to a trivialization of $\det(\tilde{L}_\gamma)$ in a way which respects concatenation of paths. It follows that an orientation data system in Coulomb gauge gives rise to a specialized coherent orientation on the finite dimensional approximation $(B(2R) \cap W^\lambda)^\sigma / S^1$. Once we fix these orientations, we obtain the following.

Proposition 13.6. *The bijective correspondences described in Propositions 13.1, 13.4 and 13.5 are orientation-preserving.*

Proof. From the above construction, we see that, on the moduli spaces of approximate Seiberg-Witten trajectories, the signs defined using P_γ^{gC} (with \mathfrak{q}^λ) coincide with those defined using \tilde{L}_γ .

Further, as λ varies on an interval $[\lambda_0, \infty]$ (including infinity), the moduli spaces $\check{M}^{\text{agC}}([x_\lambda], [y_\lambda])$ form a continuous family, carrying determinant line bundles given by $\det(P_\gamma^{\text{gC}})$. These bundles vary continuously, so the trivialization at any finite λ agrees with the one at ∞ , under the correspondence from Proposition 13.1.

The same argument applies to the other two correspondences. \square

14. THE EQUIVALENCE OF THE HOMOLOGY THEORIES

We are now ready to prove some of the results from the Introduction.

Proof of Theorem 1.1. Recall that the goal is to establish an absolutely graded isomorphism between $\tilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s}))$ and $\widetilde{HM}(Y, \mathfrak{s})$. First, by Proposition 6.6, we have that

$$(271) \quad \tilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s})) \cong \tilde{H}_*^{S^1}(\text{SWF}_{\mathfrak{q}}(Y, \mathfrak{s})),$$

where $\text{SWF}_{\mathfrak{q}}(Y, \mathfrak{s})$, defined in Section 6.1, is the analogous construction of the Floer spectrum using $\mathcal{X}_{\mathfrak{q}}^{\text{gC}} = l + p^\lambda c_{\mathfrak{q}}$ instead of $l + p^\lambda c$. Here, we require that \mathfrak{q} is a very tame, admissible perturbation satisfying the conclusions of Proposition 7.11, Proposition 8.1, and Proposition 10.2. Recall that

$$\text{SWF}_{\mathfrak{q}}(Y, \mathfrak{s}) = \Sigma^{-n(Y, \mathfrak{s}, g)} \mathbb{C} \Sigma^{-W^{(-\lambda, 0)}} I_{\mathfrak{q}}^\lambda,$$

where $I_{\mathfrak{q}}^\lambda$ is the equivariant Conley index constructed from the flow of $l + p^\lambda c_{\mathfrak{q}}$, analogous to I^λ constructed in Section 3.5.

In Propositions 8.1 and 10.2 we showed that for $\lambda = \lambda_i^\bullet \gg 0$, $\mathcal{X}_{\mathfrak{q}}^{\text{gC}}$ is a Morse-Smale equivariant quasi-gradient on $W^\lambda \cap B(2R)$. Thus, we can construct a Morse complex $(\check{C}_\lambda, \check{\partial}_\lambda)$ for $\mathcal{X}_{\mathfrak{q}}^{\text{agC}, \sigma}$ on $(W^\lambda \cap B(2R))^\sigma / S^1$ as defined in Sections 2.5 and Section 2.6. By (22), we have that

$$(272) \quad \tilde{H}_i^{S^1}(I_{\mathfrak{q}}^\lambda) \cong H_i(\check{C}_\lambda, \check{\partial}_\lambda),$$

where $0 \leq i \leq n_\lambda - 1$, where n_λ is the connectivity of the pair $(I_{\mathfrak{q}}^\lambda, (I_{\mathfrak{q}}^\lambda - (I_{\mathfrak{q}}^\lambda)^{S^1}) \cup *)$. Note that this isomorphism respects the $\mathbb{Z}[U]$ -module structures on each of the homologies.

By shifting gradings, we can rephrase (272) as

$$(273) \quad \tilde{H}_i^{S^1}(\text{SWF}_{\mathfrak{q}}(Y, \mathfrak{s})) \cong H_{i + \dim W^{(-\lambda, 0)} + 2n(Y, \mathfrak{s}, g)}(\check{C}_\lambda, \check{\partial}_\lambda),$$

for $-\dim W^{(-\lambda, 0)} - 2n(Y, \mathfrak{s}, g) \leq i \leq n_\lambda - 1 - \dim W^{(-\lambda, 0)} - 2n(Y, \mathfrak{s}, g)$. Set

$$M_\lambda := \min \{ \dim W^{(-\lambda, 0)} + 2n(Y, \mathfrak{s}, g), n_\lambda - 1 - \dim W^{(-\lambda, 0)} - 2n(Y, \mathfrak{s}, g) \},$$

so that the isomorphism in (273) holds in the grading range $[-M_\lambda, M_\lambda]$.

We claim that $M_\lambda \rightarrow \infty$ as $\lambda = \lambda_i^\bullet \rightarrow \infty$. It is clear that $\dim W^{(-\lambda, 0)} \rightarrow \infty$, so what we need to check is that

$$(274) \quad n_\lambda - \dim W^{(-\lambda, 0)} \rightarrow \infty.$$

Let us investigate how the connectivity of the pair $(I_{\mathfrak{q}}^\lambda, (I_{\mathfrak{q}}^\lambda - (I_{\mathfrak{q}}^\lambda)^{S^1}) \cup *)$ changes as we increase λ . Fix a sufficiently large eigenvalue cut-off $\mu = \lambda_i^\bullet$. (Recall that this implies that neither μ nor $-\mu$ are eigenvalues.) For $\lambda = \lambda_j^\bullet \geq \mu$, it is proved in [Man03, Section 7] (for $\mathfrak{q} = 0$, but the same argument works for any \mathfrak{q}) that

$$(275) \quad I_{\mathfrak{q}}^\lambda \simeq I_{\mathfrak{q}}^\mu \wedge I(l)^{(\mu, \lambda)}.$$

Here, $I(l)^{(\mu, \lambda)}$ is the Conley index associated to the isolated invariant set $\{0\}$ in the linear flow induced by l on the complementary subspace to W^μ in W^λ , that is, on $W^{(-\lambda, -\mu)} \oplus W^{(\mu, \lambda)}$. Let us decompose this complementary subspace according to the sign of the eigenvalues, and also according to the type of eigenvectors (connections or spinors). Specifically,

let

$$\begin{aligned} a_+^{\mu,\lambda} &= \dim(W^{(\mu,\lambda)} \cap \ker d^*), & b_+^{\mu,\lambda} &= \dim(W^{(\mu,\lambda)} \cap \Gamma(\mathbb{S})) \\ a_-^{\mu,\lambda} &= \dim(W^{(-\lambda,-\mu)} \cap \ker d^*), & b_-^{\mu,\lambda} &= \dim(W^{(-\lambda,-\mu)} \cap \Gamma(\mathbb{S})), \end{aligned}$$

where \dim denotes the real dimension. Then, the Conley index of the linear flow is

$$I(l)^{(\mu,\lambda)} \simeq D(\mathbb{R}^{a_+^{\mu,\lambda}})_+ \wedge D(\mathbb{C}^{b_+^{\mu,\lambda}})_+ \wedge (\mathbb{R}^{a_-^{\mu,\lambda}})^+ \wedge (\mathbb{C}^{b_-^{\mu,\lambda}})^+.$$

We used here the standard notation in homotopy theory: If V is a vector space, then $D(V)_+$ is the union of the unit disk in V with one disjoint basepoint, and V^+ is the one-point compactification of V .

Using (275), and keeping track of the fixed point sets in each Conley index, we obtain

$$(I_q^\lambda, (I_q^\lambda)^{S^1}) \simeq \left(D(\mathbb{R}^{a_+^{\mu,\lambda}})_+ \wedge D(\mathbb{C}^{b_+^{\mu,\lambda}})_+ \wedge \Sigma^{a_-^{\mu,\lambda}\mathbb{R}} \Sigma^{b_-^{\mu,\lambda}\mathbb{C}}(I_q^\mu), D(\mathbb{R}^{a_+^{\mu,\lambda}})_+ \wedge \Sigma^{a_-^{\mu,\lambda}\mathbb{R}}(I_q^\mu)^{S^1} \right).$$

Observe that when we change a pair (X, Y) into $(D(\mathbb{R})_+ \wedge X, D(\mathbb{R})_+ \wedge Y)$, $(D(\mathbb{C})_+ \wedge X, Y)$, $(\Sigma^{\mathbb{R}} X, \Sigma^{\mathbb{R}} Y)$, or $(\Sigma^{\mathbb{C}} X, Y)$, then the connectivity of the pair $(X, (X - Y) \cup *)$ increases by 0, 2, 1 and 2, respectively. Therefore,

$$n_\lambda = n_\mu + 2b_+^{\mu,\lambda} + a_-^{\mu,\lambda} + 2b_-^{\mu,\lambda}.$$

Since $\dim W^{(-\lambda,0)} = \dim W^{(-\mu,0)} + a_-^{\mu,\lambda} + 2b_-^{\mu,\lambda}$, we see that

$$n_\lambda - \dim W^{(-\lambda,0)} = n_\mu - \dim W^{(-\mu,0)} + 2b_+^{\mu,\lambda} \rightarrow \infty \text{ as } \lambda = \lambda_i^\bullet \rightarrow \infty.$$

This proves (274), and we conclude that $M_\lambda \rightarrow \infty$. Thus, the isomorphism (273) holds in a grading range that gets larger as $\lambda = \lambda_i^\bullet \rightarrow \infty$.

While the chain groups of \check{C}_λ consist of stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$, the natural gradings from Morse homology are not the absolute gradings on stationary points that we have been working with throughout this paper. Recall that the absolute grading $\text{gr}_\lambda^{\text{SWF}}$ on stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ defined in (231) includes a shift by $\dim W^{(-\lambda,0)} + 2n(Y, \mathfrak{s}, g)$. Therefore, we

define the complex $\widetilde{\mathcal{CM}}^\lambda(Y, \mathfrak{s}, q)$ by taking the Morse complex $(\check{C}_\lambda, \check{\partial}_\lambda)$ and shifting gradings down by $-\dim W^{(-\lambda,0)} + 2n(Y, \mathfrak{s}, g)$, so stationary points have grading given by $\text{gr}_\lambda^{\text{SWF}}$. We denote the homology of $\widetilde{\mathcal{CM}}^\lambda$ by \widetilde{HM}^λ . With this definition, we have

$$(276) \quad \widetilde{H}_i^{S^1}(\text{SWF}_q(Y, \mathfrak{s})) \cong \widetilde{HM}_i^\lambda(Y, \mathfrak{s}, q)$$

for $i \in [-M_\lambda, M_\lambda]$. Therefore, it remains to establish an isomorphism between $\widetilde{HM}_i^\lambda(Y, \mathfrak{s}, q)$ and $\widetilde{HM}_i(Y, \mathfrak{s}, q)$. In other words, we must relate the Morse homology defined in terms of stationary points and trajectories of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ versus the Floer homology defined in terms of stationary points and trajectories of \mathcal{X}_q^σ .

First, by the work of Section 5 (more precisely described in Section 5.14), we can compute $\widetilde{\mathcal{CM}}(Y, \mathfrak{s}, q)$ using stationary points and trajectories of $\mathcal{X}_q^{\text{agC},\sigma}$ instead of \mathcal{X}_q^σ . In a fixed grading range $[-N, N]$, Proposition 9.11 shows that for $\lambda = \lambda_i^\bullet \gg 0$, there is a grading-preserving identification between the stationary points of $\mathcal{X}_q^{\text{agC},\sigma}$ and the stationary points of $\mathcal{X}_{q^\lambda}^{\text{agC},\sigma}$ in $(W^\lambda \cap B(2R))^\sigma / S^1$, and thus we have an identification between the chain groups $\widetilde{\mathcal{CM}}^\lambda(Y, \mathfrak{s}, q)$ and $\widetilde{\mathcal{CM}}(Y, \mathfrak{s}, q)$ in the grading range $[-N, N]$.

By the work of Section 13, we have established an identification between the signed counts of index one (respectively boundary-obstructed index two) trajectories for $\mathcal{X}_q^{\text{agC},\sigma}$ and for

$\mathcal{X}_{q^\lambda}^{\text{agC}, \sigma}$, in the given grading range. Since these are the moduli spaces that are counted in the differentials for $\widetilde{\mathcal{CM}}^\lambda$ and $\widetilde{\mathcal{CM}}$ (defined in (11) and (54) respectively), we obtain a chain complex isomorphism in this grading range. Therefore, we have that

$$(277) \quad \widetilde{HM}_i^\lambda(Y, \mathfrak{s}) \cong \widetilde{HM}_i(Y, \mathfrak{s})$$

for $i \in [-N+1, N-1]$.

By Proposition 13.5, the U -maps on these chain complexes, defined in (19) and (57) for $\widetilde{\mathcal{CM}}^\lambda(Y, \mathfrak{s}, \mathfrak{q})$ and $\widetilde{\mathcal{CM}}(Y, \mathfrak{s}, \mathfrak{q})$ respectively, must agree in the given grading range. Thus, the isomorphism in (277) respects the U -action.

Therefore, we have established an isomorphism between $\widetilde{HM}_i^\lambda(Y, \mathfrak{s})$ and $\widetilde{HM}_i(Y, \mathfrak{s})$ $\mathbb{Z}[U]$ -modules for $i \in [-N+1, N-1]$. Take $\lambda \gg 0$ so that $M_\lambda \geq N-1$. Combining this last isomorphism with (271) and (276), we see that

$$\widetilde{H}_i^{S^1}(\text{SWF}(Y, \mathfrak{s})) \cong \widetilde{HM}_i(Y, \mathfrak{s}),$$

for $i \in [-N+1, N-1]$. Again, this isomorphism respects the U -action. Since N was arbitrary, we obtain the desired result. \square

Proof of Corollary 1.2. Recall that Bloom's variant of monopole Floer homology, $\widetilde{HM}(Y, \mathfrak{s})$, is defined to be the homology of the mapping cone (with some grading shifts) of the U -map on $\widetilde{\mathcal{CM}}(Y, \mathfrak{s}, \mathfrak{q})$.

If F is a field, the long exact sequence induced by the U -map on homology shows that $\widetilde{HM}(Y, \mathfrak{s}; F)$ is isomorphic to the homology of the mapping cone of the U -map on $\widetilde{HM}(Y, \mathfrak{s}; F)$, thought of as a complex equipped with the trivial differential. By Theorem 1.1, $\widetilde{HM}(Y, \mathfrak{s}; F)$ is calculated as the mapping cone of the U -map on $\widetilde{H}_*^{S^1}(\text{SWF}(Y, \mathfrak{s}); F)$. Applying the following lemma with $\mathbb{Z}/p^k\mathbb{Z}$ - and \mathbb{Q} -coefficients, together with the universal coefficient theorem, completes the proof. \square

Lemma 14.1. *Let X be a based, finite S^1 -CW complex. Then, for any field F , we have a graded isomorphism*

$$\widetilde{H}^n(X; F) \cong \widetilde{H}^n([H_{S^1}^*(X; F)]_1 \xrightarrow{U} \widetilde{H}_{S^1}^{*-2}(X; F)).$$

Here, $[C]_1$ means we are shifting the degree of C by 1.

Proof. Since the diagonal action of S^1 is free on $X \wedge ES_+^1$, we have that the orbit map $X \wedge ES_+^1 \rightarrow X \wedge_{S^1} ES_+^1 := (X \wedge ES_+^1)/S^1$ is a principal S^1 -bundle (away from the basepoint). This bundle is isomorphic to the pullback bundle of the map $\pi : X \wedge_{S^1} ES_+^1 \rightarrow BS^1$. There is an associated Gysin sequence:

$$\dots \rightarrow \widetilde{H}^n(X \wedge ES_+^1; F) \rightarrow \widetilde{H}^{n-1}(X \wedge_{S^1} ES_+^1; F) \xrightarrow{e \cup} \widetilde{H}^{n+1}(X \wedge_{S^1} ES_+^1; F) \rightarrow \dots,$$

where e is the Euler class of this S^1 -bundle. Since the Euler class agrees with the first Chern class, e is by definition $\pi^*(U)$, where U is the generator of $H^2(BS^1; F)$. Because ES^1 is contractible, we can rewrite the Gysin sequence as

$$\dots \rightarrow \widetilde{H}^n(X; F) \rightarrow \widetilde{H}_{S^1}^{n-1}(X; F) \xrightarrow{\pi^*(U) \cup} \widetilde{H}_{S^1}^{n+1}(X; F) \rightarrow \dots$$

However, cup product with $\pi^*(U)$ corresponds precisely to the U -action in the $F[U]$ -module structure on $H_{S^1}^*(X; \mathbb{Z})$. The Gysin sequence now gives the desired isomorphism. \square

Proof of Corollary 1.3. This follows from Theorem 1.1, given that δ and $-h$ are obtained in the same way from \widetilde{HM} resp. $\widetilde{H}_*^{S^1}(\text{SWF})$; they are both half of the grading of the lowest nontrivial element in the infinite U -tail of the corresponding homology group with \mathbb{Q} coefficients. \square

Proof of Corollary 1.4. By the Seiberg-Witten / Heegaard Floer equivalence [KLT10a, CGH12b, CGH12a], the minus and hat flavors of Heegaard Floer homology are isomorphic to \widetilde{HM} resp. \widehat{HM} . Thus, the first two isomorphisms in Corollary 1.4 are consequences of Theorem 1.1 and Corollary 1.2.

To get the isomorphism between HF^- and co-Borel homology, we can apply the isomorphism between HF^+ and Borel homology to Y with the orientation reversed. Indeed, there is a duality between $HF^-(-Y, \mathfrak{s})$ and the cohomology version of $HF^+(Y, \mathfrak{s})$; see [OS04, Proposition 2.5]. Similarly, co-Borel homology is isomorphic (up to a reversal of degrees) to the Borel cohomology of the dual spectrum, and $\text{SWF}(Y, \mathfrak{s})$ is dual to $\text{SWF}(-Y, \mathfrak{s})$ by [Man03, §9, Remark 2]. Cohomology theories can be related back to homologies by the universal coefficient theorem.

To get the isomorphism between HF^∞ and Tate homology, note that HF^∞ can be obtained from HF^- by inverting the variable U . Similarly, Tate homology is obtained from co-Borel homology by inverting U . \square

REFERENCES

- [AD14] Michèle Audin and Mihai Damian. *Morse theory and Floer homology*. Universitext. Springer, London; EDP Sciences, Les Ulis, 2014. Translated from the 2010 French original by Reinie Ern .
- [BH04] Augustin Banyaga and David Hurtubise. *Lectures on Morse homology*, volume 29 of *Kluwer Texts in the Mathematical Sciences*. Kluwer Academic Publishers Group, Dordrecht, 2004.
- [Blo11] Jonathan M. Bloom. A link surgery spectral sequence in monopole Floer homology. *Adv. Math.*, 226(4):3216–3281, 2011.
- [BNR12] Maciej Borodzik, Andr s N methi, and Andrew Ranicki. Morse theory for manifolds with boundary. Preprint, [arXiv:1207.3066](#), 2012.
- [Bot88] Raoul Bott. Morse theory indomitable. *Inst. Hautes  tudes Sci. Publ. Math.*, (68):99–114 (1989), 1988.
- [CGH12a] Vincent Colin, Paolo Ghiggini, and Ko Honda. The equivalence of Heegaard Floer homology and embedded contact homology III: from hat to plus. Preprint, [arXiv:1208.1526](#), 2012.
- [CGH12b] Vincent Colin, Paolo Ghiggini, and Ko Honda. The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions I. Preprint, [arXiv:1208.1074](#), 2012.
- [CGH12c] Vincent Colin, Paolo Ghiggini, and Ko Honda. The equivalence of Heegaard Floer homology and embedded contact homology via open book decompositions II. Preprint, [arXiv:1208.1077](#), 2012.
- [Con78] Charles Conley. *Isolated invariant sets and the Morse index*, volume 38 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I., 1978.
- [Exc] Mathematics Stack Exchange. Nonlinear function continuous but not bounded. <http://math.stackexchange.com/questions/29498/nonlinear-function-continuous-but-not-bounded>.
- [Flo87] Andreas Floer. A refinement of the Conley index and an application to the stability of hyperbolic invariant sets. *Ergodic Theory Dynam. Systems*, 7(1):93–103, 1987.
- [Flo89] Andreas Floer. Witten’s complex and infinite-dimensional Morse theory. *J. Differential Geom.*, 30(1):207–221, 1989.
- [Fra86] Robert Franzosa. Index filtrations and the homology index braid for partially ordered Morse decompositions. *Trans. Amer. Math. Soc.*, 298(1):193–213, 1986.
- [Fr 10] Kim A. Fr yshov. Monopole Floer homology for rational homology 3-spheres. *Duke Math. J.*, 155(3):519–576, 2010.
- [GM95] J. P. C. Greenlees and J. P. May. Generalized Tate cohomology. *Mem. Amer. Math. Soc.*, 113(543):viii+178, 1995.
- [Hel09] Juliette Hell. *Conley index at infinity*. PhD thesis, Freie Universit t Berlin, 2009. <http://dynamics.mi.fu-berlin.de/persons/hell.php>.
- [Kha13] Tirasan Khandhawit. *Twisted Manolescu-Floer Spectra for Seiberg-Witten Monopoles*. ProQuest LLC, Ann Arbor, MI, 2013. Thesis (Ph.D.)—Massachusetts Institute of Technology.
- [KLT10a] Ca atay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes. HF=HM I : Heegaard Floer homology and Seiberg–Witten Floer homology. Preprint, [arXiv:1007.1979](#), 2010.
- [KLT10b] Ca atay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes. HF=HM II : Reeb orbits and holomorphic curves for the ech/Heegaard-Floer correspondence. Preprint, [arXiv:1008.1595](#), 2010.
- [KLT10c] Ca atay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes. HF=HM III : Holomorphic curves and the differential for the ech/Heegaard Floer correspondence. Preprint, [arXiv:1010.3456](#), 2010.
- [KLT11] Ca atay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes. HF=HM IV : The Seiberg-Witten Floer homology and ech correspondence. Preprint, [arXiv:1107.2297](#), 2011.
- [KLT12] Ca atay Kutluhan, Yi-Jen Lee, and Clifford H. Taubes. HF=HM V : SeibergWitten Floer homology and handle additions. Preprint, [arXiv:1204.0115](#), 2012.
- [KM02] Peter B. Kronheimer and Ciprian Manolescu. Periodic Floer pro-spectra from the Seiberg-Witten equations. Preprint, [arXiv:math/0203243v2](#), 2002.
- [KM07] Peter B. Kronheimer and Tomasz S. Mrowka. *Monopoles and three-manifolds*, volume 10 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2007.
- [KMOS07] Peter B. Kronheimer, Tomasz S. Mrowka, Peter S. Ozsv th, and Zoltan Szab . Monopoles and lens space surgeries. *Ann. of Math. (2)*, 165(2):457–546, 2007.
- [Lau11] Fran ois Laudenbach. A Morse complex on manifolds with boundary. *Geom. Dedicata*, 153:47–57, 2011.

- [Lin15] Jianfeng Lin. Pin(2)-equivariant KO-theory and intersection forms of spin 4-manifolds. *Algebr. Geom. Topol.*, 15(2):863–902, 2015.
- [LM14] Tye Lidman and Ciprian Manolescu. Floer homology and covering spaces. Preprint, 2014.
- [Man03] Ciprian Manolescu. Seiberg-Witten-Floer stable homotopy type of three-manifolds with $b_1 = 0$. *Geom. Topol.*, 7:889–932 (electronic), 2003.
- [Man07] Ciprian Manolescu. A gluing theorem for the relative Bauer-Furuta invariants. *J. Differential Geom.*, 76(1):117–153, 2007.
- [Man14] Ciprian Manolescu. On the intersection forms of spin four-manifolds with boundary. *Math. Ann.*, 359(3-4):695–728, 2014.
- [Man16] Ciprian Manolescu. Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture. *J. Amer. Math. Soc.*, 29(1):147–176, 2016.
- [McC88] Christopher McCord. The connection map for attractor-repeller pairs. *Trans. Amer. Math. Soc.*, 307(1):195–203, 1988.
- [Mil65] John Milnor. *Lectures on the h-cobordism theorem*. Notes by L. Siebenmann and J. Sondow. Princeton University Press, Princeton, N.J., 1965.
- [MOY97] Tomasz Mrowka, Peter Ozsváth, and Baozhen Yu. Seiberg-Witten monopoles on Seifert fibered spaces. *Comm. Anal. Geom.*, 5(4):685–791, 1997.
- [MW01] Matilde Marcolli and Bai-Ling Wang. Equivariant Seiberg-Witten Floer homology. *Comm. Anal. Geom.*, 9(3):451–639, 2001.
- [OS03] Peter S. Ozsváth and Zoltán Szabó. On the Floer homology of plumbed three-manifolds. *Geom. Topol.*, 7:185–224, 2003.
- [OS04] Peter S. Ozsváth and Zoltán Szabó. Holomorphic disks and three-manifold invariants: properties and applications. *Ann. of Math. (2)*, 159(3):1159–1245, 2004.
- [OS11] Peter S. Ozsváth and Zoltán Szabó. Knot Floer homology and rational surgeries. *Algebr. Geom. Topol.*, 11(1):1–68, 2011.
- [PdM82] Jacob Palis, Jr. and Welington de Melo. *Geometric theory of dynamical systems*. Springer-Verlag, New York-Berlin, 1982. An introduction, Translated from the Portuguese by A. K. Manning.
- [Pru99] Artur M. Prusko. The Conley index for flows preserving generalized symmetries. In *Conley index theory (Warsaw, 1997)*, volume 47 of *Banach Center Publ.*, pages 193–217. Polish Acad. Sci., Warsaw, 1999.
- [Sal90] Dietmar Salamon. Morse theory, the Conley index and Floer homology. *Bull. London Math. Soc.*, 22(2):113–140, 1990.
- [Sas14] Hirofumi Sasahira. Gluing formula for the stable cohomotopy version of Seiberg-Witten invariants along 3-manifolds with $b_1 > 0$. Preprint, arXiv:1408.2623, 2014.
- [Sch93] Matthias Schwarz. *Morse homology*, volume 111 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1993.
- [Sto] Matthew Stoffregen. Pin(2)-equivariant Seiberg-Witten Floer homology of Seifert fibrations. Preprint, arXiv:1505.03234.
- [SW94a] Nathan Seiberg and Edward Witten. Electric-magnetic duality, monopole condensation, and confinement in $N = 2$ supersymmetric Yang-Mills theory. *Nuclear Phys. B*, 426(1):19–52, 1994.
- [SW94b] Nathan Seiberg and Edward Witten. Monopoles, duality and chiral symmetry breaking in $N = 2$ supersymmetric QCD. *Nuclear Phys. B*, 431(3):484–550, 1994.
- [Tau07] Clifford Henry Taubes. The Seiberg-Witten equations and the Weinstein conjecture. *Geom. Topol.*, 11:2117–2202, 2007.
- [Tau10] Clifford Henry Taubes. Embedded contact homology and Seiberg-Witten Floer cohomology I-V. *Geom. Topol.*, 14(5):2497–3000, 2010.
- [Wit82] Edward Witten. Supersymmetry and Morse theory. *J. Differential Geom.*, 17(4):661–692 (1983), 1982.
- [Wit94] Edward Witten. Monopoles and four-manifolds. *Math. Res. Lett.*, 1:769–796, 1994.

SCHOOL OF MATHEMATICS, INSTITUTE FOR ADVANCED STUDY, EINSTEIN DRIVE, PRINCETON, NJ 08540
E-mail address: `tlid@math.ias.edu`

DEPARTMENT OF MATHEMATICS, UCLA, 520 PORTOLA PLAZA, LOS ANGELES, CA 90095
E-mail address: `cm@math.ucla.edu`